Distribution-based Bisimulation and Bisimulation Metric in Probabilistic Automata☆

Yuan Feng[a], Lei Song[b,c], Lijun Zhang[d]

[a]*University of Technology Sydney, Australia*
[b]*Saarland University, Germany*
[c]*Max-Planck-Institut für Informatik, Germany*
[d]*State Key Laboratory of Computer Science, Institute of Software, Chinese Academy of Sciences, China*

arXiv:1512.05027v1 [cs.FL] 16 Dec 2015

Abstract

Probabilistic automata were introduced by Rabin in 1963 as language acceptors. Two automata are equivalent if and only if they accept each word with the same probability. On the other side, in the process algebra community, probabilistic automata were re-proposed by Segala in 1995 which are more general than Rabin's automata. Bisimulations have been proposed for Segala's automata to characterize the equivalence between them. So far the two notions of equivalences and their characteristics have been studied mostly independently. In this paper, we consider Segala's automata, and propose a novel notion of distribution-based bisimulation by joining the existing equivalence and bisimilarities. We demonstrate the utility of our definition by studying distribution-based bisimulation metrics, which gives rise to a robust notion of equivalence for Rabin's automata. We compare our notions of bisimulation to some existing distribution-based bisimulations and discuss their compositionality and relations to trace equivalence. Finally, we show the decidability and complexity of all relations.

Keywords: Distribution-based bisimulation, Probabilistic automota, Trace equivalence, Bisimulation metric

1. Introduction

In 1963, Rabin [1] introduced the model *probabilistic automata* as language acceptors. In a probabilistic automaton, each input symbol determines a stochastic transition matrix over the state space. Starting with the initial distribution, each word (a sequence of symbols) has a corresponding probability of reaching one of the final states, which is referred to as the accepting probability. Two automata are equivalent if and only if they accept each word with the same probability. The corresponding decision algorithm has been extensively studied, see [1, 2, 3, 4].

Markov decision processes (MDPs) were known as early as the 1950s [5], and are a popular modelling formalism used for instance in operations research, automated planning, and decision support systems. In MDPs, each state has a set of enabled actions and each enabled action leads to a distribution over successor states. MDPs have been widely used in the formal verification of randomized concurrent systems [6], and are now supported by probabilistic model checking tools such as PRISM [7], MRMC [8] and IscasMC [9].

On the other side, in the context of concurrent systems, probabilistic automata were re-proposed by Segala in 1995 [10], which extend MDPs with internal nondeterministic choices. Segala's automata are more general than Rabin's automata, in the sense that each input symbol may correspond to more than one stochastic transition matrices. Various behavioral equivalences were defined, including strong bisimulations, strong probabilistic bisimulations, and weak bisimulation extensions [10]. Strong bisimulations require all transitions being matched by equivalent states, whereas weak bisimulations allow single transition being matched by a finite execution fragment. These behavioral equivalences are used as powerful tools for state space reduction and hierarchical verification of complex systems.

☆Supported by the National Natural Science Foundation of China (NSFC) (Grant No. 61428208, 61472473 and 61361136002), Australian Research Council (ARC) under grant No DP130102764, and by the Overseas Team Program of Academy of Mathematics and Systems Science, CAS and the CAS/SAFEA International Partnership Program for Creative Research Team.

Thus, their decision algorithms [11, 12, 13] and logical characterizations [14, 15, 16] were widely studied in the literature.

Equivalences are defined for the specific initial distributions over Rabin's automata which are *deterministic* with respect to the input word. On the other side, bisimulations are usually defined over states over Segala's automata which are *non-deterministic*, i.e., an input word induces often more than one probability distributions. For Segala's automata, state-based bisimulations have arguably too strong distinguishing power, thus various relaxations have been proposed, recently. In [17], a distribution-based bisimulation was defined for Rabin's automata. Because of the deterministic nature, they proved further that this turns out to be an equivalent characterization of the equivalence in a coinductive manner, as for bisimulations. Later, this leads to, a distribution-based weak bisimulation in [18]. Since states can be matched by distributions, this novel equivalence relation ignores the branching structure of the model, and produces weaker relations over states than classical weak bisimulations [10].

The induced distribution-based strong bisimulation was further studied in [19]. Interestingly, whereas the weak bisimulation is weaker, it was shown that the distribution-based strong bisimulation agrees with the state-based bisimulations when lifted to distributions. This is rooted in the formulation of the weak bisimulation in [18]: constraints are proposed for each state in the distribution independently. This may appear too fine for the overall behaviour of the distribution: this limitation was illustrated in an example at the end of the paper [18]. Other formulations of distribution-based bisimulations have been proposed recently. In [20], weaker notions of weak bisimulations are further studied. But the corresponding strong bisimulations have not been discussed. Recently, in [21], a new distribution-based strong bisimulation is proposed for Segala's automata and extensions with continuous state space.

As one contribution of this paper, we consider Segala's probabilistic automata, and propose a novel notion of strong distribution-based bisimulation. The novel relation is coarser than the relations in [19, 18]: we show that for Rabin's probabilistic automata it coincides with equivalences, and for Segala's probabilistic automata, it is reasonably weaker than the existing bisimulation relations. Thus, it joins the two equivalence notions restricting to the corresponding sub-models.

Another contribution of this paper is the characterization of distribution-based bisimulation metrics. Bisimulations for probabilistic systems are known to be very sensitive to the transition probabilities: even a tiny perturbation of the transition probabilities will destroy bisimilarity. Thus, bisimulation metrics have been proposed [22]: the distance between any two states are measured, and the smaller the distance is, the more similar they are. If the distance is zero, one then has the classical bisimulation. Because of the nice property of robustness, bisimulation metrics have attracted a lot attentions on MDPs and their extensions with continuous state space, see [23, 24, 25, 26, 27, 28, 29, 30, 31, 32]. All of the existing bisimulation metrics mentioned above are state-based. On the other side, as states lead to distributions in MDPs, the metrics must be lifted to distributions. In the second part of the paper, we propose a distribution-based bisimulation metric; we consider it being more natural as no lifting of distances is needed. We provide a coinductive definition as well as a fixed point characterization, both of which are used in defining the state-based bisimulation metrics in the literature. We provide a logical characterization for this metric as well, and discuss the relation of our definition and the state-based ones. A direct byproduct of our bisimulation-based metrics is the notion of equivalence metric for Rabin's probabilistic automata. As for bisimulation metrics, the equivalence metric provides a robust solution for comparing Rabin's automata. To the best of our knowledge, this has not been studied in the literature.

Lastly, we consider the strong bisimulation obtained from [20], and investigate the relations of it, together with our relation and the one defined in [21]. We show that our distribution-based bisimulation is the coarsest among them. In more detail, we show that even though all the distribution-based bisimulations are not compositional in general, they are compositional if restricted to a subclass of schedulers. We compare all notions of bisimulation to trace equivalences, and show that they are all finer than a priori trace distribution equivalence, but incomparable to trace equivalences. Further, we study the corresponding decision algorithms in this paper. In [21], the authors have studied the decision algorithm for their bisimulation, and they pointed out that it can be used for deciding other variants as well. In contrast to state-based bisimulation metrics, the problem of computing distribution-based bisimulation metrics is much harder. Actually, we prove that the problem of computing the distribution-based bisimulation metrics without discounting is undecidable. However, if equipped with a discounting factor, the problem turns to be decidable and is NP-hard.

This paper is an extended version of the conference paper [33]. In addition to the conference version, we have added missing proofs, investigated the relationship of our bisimulation with other distribution-based bisimulations

in the literature, and shown their compositionality with respect to a subclass of schedulers. We also discussed the decidability and complexity of computing all mentioned relations.

Organization of the paper. We discuss related works in Section 2. We introduce some notations in Section 3. Section 4 recalls the definitions of probabilistic automata, equivalence, and bisimulation relations. We present our distribution-based bisimulation in Section 5, and bisimulation metrics and their logical characterizations in Section 6. In Section 7 we recall some existing notions of bisimulation in the literature and compare them with our bisimulation. We also discuss their compositionality and relations to trace equivalences. The decidability and complexity of distribution-based approximate bisimulation are presented in Section 8, while Section 9 concludes the paper.

2. Related Works

In probabilistic verification, bisimulation-based behavioral equivalences are often used in abstracting the original system by aggregating bisimilar states together. Coarser bisimulation thus leads to smaller quotient system through the aggregation process. On the other side, smaller quotient may lose properties of the original system. The logical characterization problem studies the relationship between *bisimilar states* and *logical equivalent states.*

For Segala's automata, he has already investigated the relationship of behavioural equivalences and logical equivalences with respect to the logic PCTL (probabilistic computational tree logic). It was shown that strong bisimulation preserves PCTL properties, and weak bisimulation preserves a PCTL fragment without the next operator [10, 34]. Moreover, these bisimulations are strictly finer than PCTL equivalence, i.e., they distinguish even states which satisfy the same set of PCTL formulas. In [35], a novel coarser bisimulation was proposed which agrees with the PCTL logical equivalences. Extensions of the Hennessy-Milner logic of Larsen and Skou [36] were also extensively studied in the literature in this respect, including [37, 14, 15, 38, 39].

PCTL logical formulas have atomic propositions to characterize state properties, and can express more involved nested properties. For a simple class of properties, such as the probabilistic reachability, the *state-based* bisimulations are arguably too fine grained. In the literature several authors proposed preorders based on asymmetric simulation relations [34, 40]. Recently, this has led to further development of several distribution-based symmetric bisimulations [18, 19, 20, 33, 21], as discussed in the introduction.

To construct the quotient system with respect to a bisimulation, one needs to decide whether two states or distributions are bisimilar. Thus, decision algorithm for bisimulations is a fundamental problem, and has been extensively studied in the literature. This rooted in the *partition refinement algorithm* for the classical transition system. For Segala's automata, while state-based bisimulation can be decided in polynomial time [12, 11, 41, 42, 13], decision procedures for distribution-based bisimulation are more expensive than the ones for state-based bisimulation [43, 44].

3. Preliminaries

Distributions.. For a finite set S, a (probability) distribution is a function $\mu : S \to [0,1]$ satisfying $|\mu| := \sum_{s\in S} \mu(s) = 1$. We denote by $Dist(S)$ the set of distributions over S. We shall use $s, r, t, \ldots$ and $\mu, \nu \ldots$ to range over S and $Dist(S)$, respectively. Given a set of distributions $\{\mu_i\}_{1\le i\le n}$, and a set of positive weights $\{p_i\}_{1\le i\le n}$ such that $\sum_{1\le i\le n} p_i = 1$, the *convex combination* $\mu = \sum_{1\le i\le n} p_i \cdot \mu_i$ is the distribution such that $\mu(s) = \sum_{1\le i\le n} p_i \cdot \mu_i(s)$ for each $s \in S$. The support of μ is defined by $supp(\mu) := \{s \in S \mid \mu(s) > 0\}$. For an equivalence relation R defined on S, we write $\mu R \nu$ if it holds that $\mu(C) = \nu(C)$ for all equivalence classes $C \in S/R$. A distribution μ is called *Dirac* if $|supp(\mu)| = 1$, and we let δ_s denote the Dirac distribution with $\delta_s(s) = 1$.

Note that when S is finite and an order over S is fixed, the distributions $Dist(S)$ over S, when regarded as a subset of $\mathbb{R}^{|S|}$, is both convex and compact. In this paper, when we talk about convergence of distributions, or continuity of relations such as transitions, bisimulations, and pseudo-metrics between distributions, we are referring to the normal topology of $\mathbb{R}^{|S|}$. For a set $F \subseteq S$, we define the characteristic (column) vector η_F by letting $\eta_F(s) = 1$ if $s \in F$, and 0 otherwise.

pseudo-metric.. A pseudo-metric over $Dist(S)$ is a function $d : Dist(S) \times Dist(S) \to [0,1]$ such that (i) $d(\mu,\mu) = 0$; (ii) $d(\mu,\nu) = d(\nu,\mu)$; (iii) $d(\mu,\nu) + d(\nu,\omega) \ge d(\mu,\omega)$. In this paper, we assume that a pseudo-metric is continuous.

4. Probabilistic Automata and Bisimulations

4.1. Probabilistic Automata

Let AP be a finite set of atomic propositions. We recall the notion of probabilistic automata introduced by Segala [10].

Definition 4.1 (Probabilistic Automata). *A* probabilistic automaton *is a tuple $\mathcal{A} = (S, Act, \rightarrow, L, \alpha)$ where S is a finite set of states, Act is a finite set of actions, $\rightarrow \subseteq S \times Act \times Dist(S)$ is a transition relation, $L : S \rightarrow 2^{AP}$ is a labelling function, and $\alpha \in Dist(S)$ is the initial distribution.*

As usual we only consider image-finite probabilistic automata, i.e. for all $s \in S$, the set $\{\mu \mid (s, a, \mu) \in \rightarrow\}$ is finite. A transition $(s, a, \mu) \in \rightarrow$ is denoted by $s \xrightarrow{a} \mu$. We denote by $EA(s) := \{a \mid s \xrightarrow{a} \mu\}$ the set of enabled actions in s. We say $\mathcal{A}$ is *input-enabled*, if $EA(s) = Act$ for all $s \in S$. The state s is *deterministic* if $(s, a, \mu) \in \rightarrow$ and $(s, a, \mu') \in \rightarrow$ imply that $\mu = \mu'$. We say $\mathcal{A}$ is an MDP if all states in S are deterministic.

Interestingly, a subclass of probabilistic automata were already introduced by Rabin in 1963 [1]; Rabin's probabilistic automata were referred to as *reactive automata* in [10]. We adopt this convention in this paper.

Definition 4.2 (Reactive Automata). *We say $\mathcal{A}$ is* reactive *if it is input-enabled and deterministic, and for all s, $L(s) \in \{\emptyset, AP\}$.*

Here the condition $L(s) \in \{\emptyset, AP\}$ implies that the states can be partitioned into two equivalence classes according to their labelling. Below we shall identify $F := \{s \mid L(s) = AP\}$ as the set of *accepting states*, a terminology used in automata theory. In a reactive automaton, each action $a \in Act$ is enabled precisely once for all $s \in S$, thus inducing a stochastic matrix $M(a)$ satisfying $s \xrightarrow{a} M(a)(s, \cdot)$.

4.2. Probabilistic Bisimulation and Equivalence

First, we recall the definition of (strong) probabilistic bisimulation for probabilistic automata [10]. Let $\{s \xrightarrow{a} \mu_i\}_{i \in I}$ be a collection of transitions, and let $\{p_i\}_{i \in I}$ be a collection of probabilities with $\sum_{i \in I} p_i = 1$. Then $(s, a, \sum_{i \in I} p_i \cdot \mu_i)$ is called a *combined transition* and is denoted by $s \xrightarrow{a}_{\mathrm{P}} \mu$ where $\mu = \sum_{i \in I} p_i \cdot \mu_i$.

Definition 4.3 (Probabilistic bisimulation [10]). *An equivalence relation $R \subseteq S \times S$ is a probabilistic bisimulation if sRr implies that $L(s) = L(r)$, and for each $s \xrightarrow{a} \mu$, there exists a combined transition $r \xrightarrow{a}_{\mathrm{P}} \nu$ such that $\mu R \nu$.*

We write $s \sim_{\mathrm{P}} r$ whenever there is a probabilistic bisimulation R such that sRr.

Recently, in [18], a distribution-based weak bisimulation has been proposed, and the induced distribution-based strong bisimulation is further studied in [19]. Their bisimilarity is shown to be the same as $\sim_{\mathrm{P}}$ when lifted to distributions. Below we recall the definition of equivalence for reactive automata introduced by Rabin [1].

Definition 4.4 (Equivalence for Reactive Automata [1]). *Let $\mathcal{A}_i = (S_i, Act, \rightarrow_i, L_i, \alpha_i)$, $i = 1, 2$, be two reactive automata with F_i being the set of final states for $\mathcal{A}_i$. We say $\mathcal{A}_1$ and $\mathcal{A}_2$ are equivalent if $\mathcal{A}_1(w) = \mathcal{A}_2(w)$ for each $w \in Act^*$, where $\mathcal{A}_i(w) := \alpha_i M_i(a_1) \ldots M_i(a_k) \eta_{F_i}$ provided $w = a_1 \ldots a_k$.*

Stated in plain english, $\mathcal{A}_1$ and $\mathcal{A}_2$ with the same set of actions are equivalent iff for an arbitrary input w, $\mathcal{A}_1$ and $\mathcal{A}_2$ accept w with the same probability.

So far bisimulations and equivalences were studied mostly independently. The only exception we are aware is [17], in which for Rabin's probabilistic automata, a distribution-based bisimulation is defined that generalizes both equivalence and bisimulations.

Definition 4.5 (Bisimulation for Reactive Automata [17]). *Let $\mathcal{A}_i = (S_i, Act, \rightarrow_i, L_i, \alpha_i)$, $i = 1, 2$, be two reactive automata, and F_i the set of final states for $\mathcal{A}_i$. A relation $R \subseteq Dist(S_1) \times Dist(S_2)$ is a bisimulation if for each $\mu R \nu$ it holds (i) $\mu \cdot \eta_{F_1} = \nu \cdot \eta_{F_2}$, and (ii) $(\mu M_1(a)) R (\nu M_2(a))$ for all $a \in Act$.*

We write $\mu \sim_d \nu$ whenever there is a bisimulation R such that $\mu R \nu$.

It was shown in [17] that two reactive automata are equivalent if and only if their initial distributions are distribution-based bisimilar according to the definition above.

5. A Novel Bisimulation Relation

In this section we introduce a notion of distribution-based bisimulation for Segala's automata by extending the bisimulation defined in [17]. We shall show the compatibility of our definition with previous ones in Subsection 5.1, and some properties of our bisimulation in Subsection 5.2.

For the first step of defining a distribution-based bisimulation, we need to extend the transitions starting from states to those starting from distributions. A natural candidate for such an extension is as follows: for a distribution μ to perform an action a, each state in its support *must* make a combined a-move. However, this definition is problematic, as in Segala's general probabilistic automata, action a may not always be enabled in all support states of μ. In this paper, we deal with this problem by first defining distribution-based bisimulations (resp. distances) for input-enabled automata, for which the transition between distributions can be naturally defined, and then reducing the equivalence (resp. distance) of two distributions in a general probabilistic automaton to the bisimilarity (resp. distance) of these distributions in an input-enabled automaton which is obtained from the original one by adding a *dead* state and some additional transitions to the dead state.

To make our idea more rigorous, we need some notations. For $A \subseteq AP$ and a distribution μ, we define $\mu(A) := \sum\{\mu(s) \mid L(s) = A\}$, which is the probability of being in those state s with label A.

Definition 5.1. *We write $\mu \xrightarrow{a} \mu'$ if for each $s \in supp(\mu)$ there exists $s \xrightarrow{a}_{\mathrm{P}} \mu_s$ such that $\mu' = \sum_s \mu(s) \cdot \mu_s$.*

We first present our distribution-based bisimulation for input-enabled probabilistic automata.

Definition 5.2. *Let $\mathcal{A} = (S, Act, \rightarrow, L, \alpha)$ be an input-enabled probabilistic automaton. A symmetric relation $R \subseteq Dist(S) \times Dist(S)$ is a (distribution-based) bisimulation if $\mu R \nu$ implies that*

1. *$\mu(A) = \nu(A)$ for each $A \subseteq AP$, and*
2. *for each $a \in Act$, whenever $\mu \xrightarrow{a} \mu'$ then there exists a transition $\nu \xrightarrow{a} \nu'$ such that $\mu' R \nu'$.*

We write $\mu \sim^{\mathcal{A}} \nu$ if there is a bisimulation R such that $\mu R \nu$.

Obviously, the bisimilarity $\sim^{\mathcal{A}}$ is the largest bisimulation relation over $Dist(S)$.

For probabilistic automata which are not input-enabled, we define distribution-based bisimulation with the help of *input-enabled extension* specified as follows.

Definition 5.3. *Let $\mathcal{A} = (S, Act, \rightarrow, L, \alpha)$ be a probabilistic automaton over AP. The* input-enabled extension *of $\mathcal{A}$, denoted by $\mathcal{A}_\perp$, is defined as an (input-enabled) probabilistic automaton $(S_\perp, Act, \rightarrow_\perp, L_\perp, \alpha)$ over $AP_\perp$ where*

1. *$S_\perp = S \cup \{\perp\}$ where $\perp$ is a* dead *state not in S;*
2. *$AP_\perp = AP \cup \{dead\}$ with $dead \notin AP$;*
3. *$\rightarrow_\perp = \rightarrow \cup \{(s, a, \delta_\perp) \mid a \notin EA(s)\} \cup \{(\perp, a, \delta_\perp) \mid a \in Act\}$;*
4. *$L_\perp(s) = L(s)$ for any $s \in S$, and $L_\perp(\perp) = \{dead\}$.*

Definition 5.4. *Let $\mathcal{A}$ be a probabilistic automaton which is not input-enabled. Then μ and ν are bisimilar, denoted by $\mu \sim^{\mathcal{A}} \nu$, if $\mu \sim^{\mathcal{A}_\perp} \nu$.*

We always omit the superscript $\mathcal{A}$ in $\sim^{\mathcal{A}}$ when no confusion arises.

5.1. Compatibility

In this subsection we instantiate appropriate labelling functions and show that our notion of bisimilarity is a conservative extension of both probabilistic bisimulation [1] and equivalence relations [17].

Lemma 5.1. *Let $\mathcal{A}$ be a probabilistic automaton where $AP = Act$, and $L(s) = EA(s)$ for each s. Then, $\mu \sim_P \nu$ implies $\mu \sim \nu$.*

PROOF. First, it is easy to see that for a given probabilistic automaton $\mathcal{A}$ with $AP = Act$ and $L(s) = EA(s)$ for each s, and distributions μ and ν in $Dist(S)$, $\mu \sim_P \nu$ in $\mathcal{A}$ if and only if $\mu \sim_P \nu$ in the input-enabled extension $\mathcal{A}_\perp$. Thus we can assume without loss of any generality that $\mathcal{A}$ itself is input-enabled.

It suffices to show that the symmetric relation

$$R = \{(\mu, \nu) \mid \mu \sim_P \nu\}$$

is a bisimulation. For each $A \subseteq Act$, let $S(A) = \{s \in S \mid L(s) = A\}$. Then $S(A)$ is the disjoint union of some equivalence classes of $\sim_P$; that is, $S(A) = \cup\{M \in S/\sim_P \mid M \cap S(A) \neq \emptyset\}$. Suppose $\mu \sim_P \nu$. Then for any $M \in S/\sim_P$, $\mu(M) = \nu(M)$, hence $\mu(A) = \mu(S(A)) = \nu(S(A)) = \nu(A)$.

Let $\mu \xrightarrow{a} \mu'$. Then for any $s \in S$ there exists $s \xrightarrow{a}_P \mu_s$ such that

$$\mu' = \sum_{s \in S} \mu(s) \cdot \mu_s.$$

Now for each $t \in S$, let $[t]_{\sim_P}$ be the equivalence class of $\sim_P$ which contains t. Then for every $s \in [t]_{\sim_P}$, to match the transition $s \xrightarrow{a}_P \mu_s$ there exists some ν_t^s such that $t \xrightarrow{a}_P \nu_t^s$ and $\mu_s \sim_P \nu_t^s$. Let

$$\nu_t = \sum_{s \in [t]_{\sim_P}} \frac{\mu(s)}{\mu([t]_{\sim_P})} \cdot \nu_t^s.$$

Then we have $t \xrightarrow{a}_P \nu_t$, and $\nu \xrightarrow{a} \nu'$ where

$$\nu' = \sum_{t \in S} \nu(t) \cdot \nu_t.$$

It remains to prove $\mu' \sim_P \nu'$. For any $M \in S/\sim_P$, since $\mu_s \sim_P \nu_t^s$ we have

$$\nu_t(M) = \sum_{s \in [t]_{\sim_P}} \frac{\mu(s)}{\mu([t]_{\sim_P})} \nu_t^s(M) = \sum_{s \in [t]_{\sim_P}} \frac{\mu(s)}{\mu([t]_{\sim_P})} \mu_s(M).$$

Thus

$$\begin{aligned} \nu'(M) &= \sum_{t \in S} \nu(t) \sum_{s \in [t]_{\sim_P}} \frac{\mu(s)}{\mu([t]_{\sim_P})} \mu_s(M) \\ &= \sum_{s \in S} \mu(s) \mu_s(M) \sum_{t \in [s]_{\sim_P}} \frac{\nu(t)}{\nu([s]_{\sim_P})} \\ &= \sum_{s \in S} \mu(s) \mu_s(M) = \mu'(M) \end{aligned}$$

where for the second equality we have used the fact that $\mu([t]_{\sim_P}) = \nu([s]_{\sim_P})$ whenever $s \sim_P t$. □

Probabilistic bisimulation is defined over distributions inside one automaton, whereas equivalence for reactive automata is defined over two automata. However, they can be connected by the notion of direct sum of two automata, which is the automaton obtained by considering the disjoint union of states, edges and labelling functions respectively.

Lemma 5.2. *Let $\mathcal{A}_1$ and $\mathcal{A}_2$ be two reactive automata with the same set of actions Act, and α_1 and α_2 the corresponding initial distributions. Then the following are equivalent:*

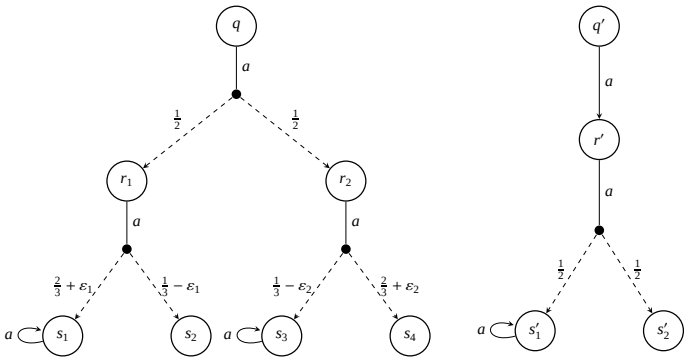

Figure 1: An illustrating example in which $L(s) = EA(s)$ for each s.

1. *$\mathcal{A}_1$ and $\mathcal{A}_2$ are equivalent,*
2. *$\alpha_1 \sim_d \alpha_2$,*
3. *$\alpha_1 \sim \alpha_2$ in their direct sum.*

PROOF. The equivalence between (1) and (2) is shown in [17]. The equivalence between (2) and (3) is straightforward, as for reactive automata our definition degenerates to Definition 4.5. □

To conclude this subsection, we present an example to show that our bisimilarity is *strictly* weaker than $\sim_P$.

Example 5.1. *Consider the example probabilistic automaton depicted in Fig. 1, which is inspired from an example in [17]. Let $AP = Act = \{a\}$, $L(s) = EA(s)$ for each s, and $\varepsilon_1 = \varepsilon_2 = 0$. We argue that $q \not\sim_P q'$. Otherwise, note $q \xrightarrow{a} \frac{1}{2}\delta_{r_1} + \frac{1}{2}\delta_{r_2}$ and $q' \xrightarrow{a} \delta_{r'}$. Then we must have $r' \sim_P r_1 \sim_P r_2$. This is impossible, as $r_1 \xrightarrow{a} \frac{2}{3}\delta_{s_1} + \frac{1}{3}\delta_{s_2}$ and $r' \xrightarrow{a} \frac{1}{2}\delta_{s'_1} + \frac{1}{2}\delta_{s'_2}$, but $s_1 \sim_P s'_1 \not\sim_P s_2 \sim_P s'_2$.*

However, by our definition of bisimulation, the Dirac distributions δ_q and $\delta_{q'}$ are indeed bisimilar. The reason is that we have the following transition

$$\frac{1}{2}\delta_{r_1} + \frac{1}{2}\delta_{r_2} \xrightarrow{a} \frac{1}{3}\delta_{s_1} + \frac{1}{6}\delta_{s_2} + \frac{1}{6}\delta_{s_3} + \frac{1}{3}\delta_{s_4},$$

and it is easy to check $\delta_{s_1} \sim \delta_{s_3} \sim \delta_{s'_1}$ and $\delta_{s_2} \sim \delta_{s_4} \sim \delta_{s'_2}$. Thus we have $\frac{1}{2}\delta_{r_1} + \frac{1}{2}\delta_{r_2} \sim \delta_{r'}$, and finally $\delta_q \sim \delta_{q'}$.

5.2. Properties of the Relations

In the following, we show that the notion of bisimilarity is in harmony with the linear combination and the limit of distributions.

Definition 5.5. *A binary relation $R \subseteq Dist(S) \times Dist(S)$ is said to be*

- linear, *if for any finite set I and any probabilistic distribution $\{p_i\}_{i\in I}$, $\mu_i R \nu_i$ for each i implies $(\sum_{i\in I} p_i \cdot \mu_i) R (\sum_{i\in I} p_i \cdot \nu_i)$;*
- continuous, *if for any convergent sequences of distributions $\{\mu_i\}_i$ and $\{\nu_i\}_i$, $\mu_i R \nu_i$ for each i implies $(\lim_i \mu_i) R (\lim_i \nu_i)$;*
- left-decomposable, *if $(\sum_{i\in I} p_i \cdot \mu_i) R \nu$, where $0 < p_i \leq 1$ and $\sum_{i\in I} p_i = 1$, then ν can be written as $\sum_{i\in I} p_i \cdot \nu_i$ such that $\mu_i R \nu_i$ for every $i \in I$.*
- left-convergent, *if $(\lim_i \mu_i) R \nu$, then for any i we have $\mu_i R \nu_i$ for some ν_i with $\lim_i \nu_i = \nu$.*

We prove below that our transition relation between distributions satisfies these properties.

Lemma 5.3. *For an input-enabled probabilistic automaton, the transition relation $\xrightarrow{a}$ between distributions is linear, continuous, left-decomposable, and left-convergent.*

PROOF.

- *Linearity*. Let I be a finite index set and $\{p_i \mid i \in I\}$ a probabilistic distribution on I. Suppose $\mu_i \xrightarrow{a} \nu_i$ for each $i \in I$. Then by definition, for each s there exists $s \xrightarrow{a}_{\mathsf{P}} \mu_s^i$ such that $\nu_i = \sum_s \mu_i(s) \cdot \mu_s^i$. Now let $\mu = \sum_{i \in I} p_i \cdot \mu_i$. Then for each $s \in supp(\mu)$,
$$s \xrightarrow{a}_{\mathsf{P}} \mu_s := \sum_{i \in I} \frac{p_i \mu_i(s)}{\mu(s)} \cdot \mu_s^i.$$
On the other hand, we check that
$$\nu := \sum_{i \in I} p_i \cdot \nu_i \quad = \quad \sum_{s \in S} \sum_{i \in I} p_i \mu_i(s) \cdot \mu_s^i = \sum_{s \in S} \mu(s) \cdot \mu_s.$$
Thus $\mu \xrightarrow{a} \nu$ as expected.

- *Continuity*. Suppose $\mu_i \xrightarrow{a} \nu_i$ for each $i \in I$, and $\lim_i \mu_i = \mu$. By definition, for each s there exists $s \xrightarrow{a}_{\mathsf{P}} \mu_s^i$ such that $\nu_i = \sum_s \mu_i(s) \cdot \mu_s^i$. Note that $Dist(S)$ is a compact set. For each s we can choose a convergent subsequence $\{\mu_s^{i_k}\}_k$ of $\{\mu_s^i\}_i$ such that $\lim_k \mu_s^{i_k} = \mu_s$ for some μ_s. Then $s \xrightarrow{a}_{\mathsf{P}} \mu_s$, and
$$\mu \xrightarrow{a} \nu := \sum_{s \in S} \mu(s) \cdot \mu_s.$$
Note that for each k,
$$\|\nu_{i_k} - \nu\|_1 \quad \leq \quad \|\mu_{i_k} - \mu\|_1 + \sum_{s \in S} \mu(s) \|\mu_s^{i_k} - \mu_s\|_1$$
where $\|\cdot\|_1$ denotes the l_1-norm. We have $\nu = \lim_k \nu_{i_k}$ by the assumption that $\lim_i \mu_i = \mu$. Thus $\lim_i \nu_i = \nu$, as $\{\nu_i\}_i$ itself converges.

- *Left-decomposability*. Let $\mu := (\sum_{i \in I} p_i \cdot \mu_i) \xrightarrow{a} \nu$. Then by definition, for each s there exists $s \xrightarrow{a}_{\mathsf{P}} \mu_s$ such that $\nu = \sum_s \mu(s) \cdot \mu_s$. Thus
$$\mu_i \xrightarrow{a} \nu_i := \sum_{s \in S} \mu_i(s) \cdot \mu_s.$$
Finally, it is easy to show that $\sum_{i \in I} p_i \cdot \nu_i = \nu$.

- *Left-convergence*. Similar to the last case. □

Theorem 5.1. *The bisimilarity relation $\sim$ is both linear and continuous.*

PROOF. Note that if $\mu_i \in Dist(S)$ for any i, then both $\sum_i p_i \cdot \mu_i$ and $\lim_i \mu_i$ (if exists) are again in $Dist(S)$. Thus we need only consider the case when the automaton is input-enabled.

- *Linearity*. It suffices to show that the symmetric relation
$$R \quad = \quad \left\{ \left(\sum_{i \in I} p_i \cdot \mu_i, \sum_{i \in I} p_i \cdot \nu_i \right) \mid I \text{ finite}, \sum_{i \in I} p_i = 1, \forall i. (p_i \geq 0 \wedge \mu_i \sim \nu_i) \right\}$$
is a bisimulation. Let $\mu = \sum_{i \in I} p_i \cdot \mu_i$, $\nu = \sum_{i \in I} p_i \cdot \nu_i$, and $\mu R \nu$. Then for any $A \subseteq AP$,
$$\mu(A) = \sum_{i \in I} p_i \cdot \mu_i(A) = \sum_{i \in I} p_i \cdot \nu_i(A) = \nu(A).$$
Now suppose $\mu \xrightarrow{a} \mu'$. Then by Lemma 5.3 (left-decomposability), for each $i \in I$ we have $\mu_i \xrightarrow{a} \mu_i'$ for some μ_i' such that $\mu' = \sum_i p_i \cdot \mu_i'$. From the assumption that $\mu_i \sim \nu_i$, we derive $\nu_i \xrightarrow{a} \nu_i'$ with $\mu_i' \sim \nu_i'$ for each i. Thus $\nu \xrightarrow{a} \nu' := \sum_i p_i \cdot \nu_i'$ by Lemma 5.3 again (linearity). Finally, it is obvious that $(\mu', \nu') \in R$.

- *Continuity*. It suffices to show that the symmetric relation

$$R = \{(\mu, \nu) \mid \forall i \geq 1, \mu_i \sim \nu_i, \lim_i \mu_i = \mu, \text{ and } \lim_i \nu_i = \nu\}$$

is a bisimulation. First, for any $A \subseteq AP$, we have

$$\mu(A) = \lim_i \mu_i(A) = \lim_i \nu_i(A) = \nu(A).$$

Let $\mu \xrightarrow{a} \mu'$. By Lemma 5.3 (left-convergence), for any i we have $\mu_i \xrightarrow{a} \mu_i'$ with $\lim_i \mu_i' = \mu'$. To match the transitions, we have $\nu_i \xrightarrow{a} \nu_i'$ such that $\mu_i' \sim \nu_i'$. Note that $Dist(S)$ is a compact set. We can choose a convergent subsequence $\{\nu_{i_k}'\}_k$ of $\{\nu_i'\}_i$ such that $\lim_k \nu_{i_k}' = \nu'$ for some ν'. From the fact that $\lim_i \nu_i = \nu$ and Lemma 5.3 (continuity), it holds $\nu \xrightarrow{a} \nu'$ as well. Finally, it is easy to see that $(\mu', \nu') \in R$.

□

In general, our definition of bisimilarity is not left-decomposable. This is in sharp contrast with the bisimulations defined by using the lifting technique [45]. However, it should not be regarded as a shortcoming; actually it is the key requirement we abandon in this paper, which makes our definition reasonably weak. This has been clearly illustrated in Example 5.1.

6. Bisimulation Metrics

We present distribution-based bisimulation metrics with discounting factor $\gamma \in (0, 1]$ in this section. Three different ways of defining bisimulation metrics between states exist in the literature: one coinductive definition based on bisimulations [46, 47, 48, 26], one based on the maximal logical differences [23, 25, 49], and one on fixed point [24, 49, 27]. We propose all the three versions for our distribution-based bisimulations with discounting. Moreover, we show that they coincide. We fix a discounting factor $\gamma \in (0, 1]$ throughout this section. For any $\mu, \nu \in Dist(S)$, we define the distance

$$d_{AP}(\mu, \nu) := \frac{1}{2} \sum_{A \subseteq AP} |\mu(A) - \nu(A)|.$$

Then it is easy to check that

$$d_{AP}(\mu, \nu) = \max_{\mathcal{B} \subseteq 2^{AP}} \left| \sum_{A \in \mathcal{B}} \mu(A) - \sum_{A \in \mathcal{B}} \nu(A) \right| = \max_{\mathcal{B} \subseteq 2^{AP}} \left[\sum_{A \in \mathcal{B}} \mu(A) - \sum_{A \in \mathcal{B}} \nu(A) \right].$$

6.1. A Direct Approach

Definition 6.1. *Let $\mathcal{A} = (S, Act, \rightarrow, L, \alpha)$ be an input-enabled probabilistic automaton. A family of symmetric relations $\{R_\varepsilon \mid \varepsilon \geq 0\}$ over $Dist(S)$ is a (discounted) approximate bisimulation if for any $\varepsilon \geq 0$ and $\mu R_\varepsilon \nu$, we have*

1. *$d_{AP}(\mu, \nu) \leq \varepsilon$;*
2. *for each $a \in Act$, $\mu \xrightarrow{a} \mu'$ implies that there exists a transition $\nu \xrightarrow{a} \nu'$ such that $\mu' R_{\varepsilon/\gamma} \nu'$.*

We write $\mu \sim_\varepsilon^{\mathcal{A}} \nu$ whenever there is an approximate bisimulation $\{R_\varepsilon \mid \varepsilon \geq 0\}$ such that $\mu R_\varepsilon \nu$. For any two distributions μ and ν, we define the bisimulation distance of μ and ν as

$$D_b^{\mathcal{A}}(\mu, \nu) = \inf\{\varepsilon \geq 0 \mid \mu \sim_\varepsilon^{\mathcal{A}} \nu\}. \tag{1}$$

Again, the approximate bisimulation and bisimulation distance of distributions in a general probabilistic automaton can be defined in terms of the corresponding notions in the input-enabled extension; that is, $\mu \sim_\varepsilon^{\mathcal{A}} \nu$ if $\mu \sim_\varepsilon^{\mathcal{A}_\perp} \nu$, and $D_b^{\mathcal{A}}(\mu, \nu) := D_b^{\mathcal{A}_\perp}(\mu, \nu)$. We always omit the superscripts for simplicity if no confusion arises.

It is standard to show that the family $\{\sim_\varepsilon | \ \varepsilon \geq 0\}$ is itself an approximate bisimulation. The following lemma collects some properties of $\sim_\varepsilon$.

Lemma 6.1. *1. For each ε, the ε-bisimilarity $\sim_\varepsilon$ is both linear and continuous.*

2. If $\mu \sim_{\varepsilon_1} \nu$ and $\nu \sim_{\varepsilon_2} \omega$, then $\mu \sim_{\varepsilon_1+\varepsilon_2} \omega$;

3. $\sim_{\varepsilon_1} \subseteq \sim_{\varepsilon_2}$ whenever $\varepsilon_1 \leq \varepsilon_2$.

PROOF. The proof of item 1 is similar to Theorem 5.1. For item 2, it suffices to show that $\{R_\varepsilon \mid \varepsilon \geq 0\}$ where $R_\varepsilon = \bigcup_{\varepsilon_1+\varepsilon_2=\varepsilon} (\sim_{\varepsilon_1} \circ \sim_{\varepsilon_2})$ is an approximate bisimulation (in the extended automaton, if necessary), which is routine. For item 3, suppose $\varepsilon_2 > 0$. It is easy to show $\{R_\varepsilon \mid \varepsilon \geq 0\}$, $R_\varepsilon = \sim_{\varepsilon\varepsilon_1/\varepsilon_2}$, is an approximate bisimulation. Then if $\mu \sim_{\varepsilon_1} \nu$, that is, $\mu \sim_{\varepsilon_2\varepsilon_1/\varepsilon_2} \nu$, we have $\mu R_{\varepsilon_2} \nu$, and thus $\mu \sim_{\varepsilon_2} \nu$ as required. □

The following theorem states that the infimum in the definition Eq. (1) of bisimulation distance can be replaced by minimum; that is, the infimum is achievable.

Theorem 6.1. *For any $\mu, \nu \in Dist(S)$, $\mu \sim_{D_b(\mu,\nu)} \nu$.*

PROOF. By definition, we need to prove $\mu \sim_{D_b(\mu,\nu)} \nu$ in the extended automaton. We first prove that for any $\varepsilon \geq 0$, the symmetric relations $\{R_\varepsilon \mid \varepsilon \geq 0\}$ where

$$R_\varepsilon = \{(\mu, \nu) \quad | \quad \mu \sim_{\varepsilon_i} \nu \text{ for each } \varepsilon_1 \geq \varepsilon_2 \geq \cdots \geq 0, \text{ and } \lim_{i\to\infty} \varepsilon_i = \varepsilon\}$$

is an approximate bisimulation. Suppose $\mu R_\varepsilon \nu$. Since $\mu \sim_{\varepsilon_i} \nu$ we have $d_{AP}(\mu, \nu) \leq \varepsilon_i$ for each i. Thus $d_{AP}(\mu, \nu) \leq \varepsilon$ as well. Furthermore, if $\mu \xrightarrow{a} \mu'$, then for any $i \geq 1$, $\nu \xrightarrow{a} \nu_i$ and $\mu' \sim_{\varepsilon_i/\gamma} \nu_i$. Since $Dist(S)$ is compact, there exists a subsequence $\{\nu_{i_k}\}_k$ of $\{\nu_i\}_i$ such that $\lim_k \nu_{i_k} = \nu'$ for some ν'. We claim that

- $\nu \xrightarrow{a} \nu'$. This follows from the continuity of the transition $\xrightarrow{a}$, Lemma 5.3.
- For each $k \geq 1$, $\mu' \sim_{\varepsilon_{i_k}/\gamma} \nu'$. Suppose conversely that $\mu' \not\sim_{\varepsilon_{i_k}/\gamma} \nu'$ for some k. Then by the continuity of $\sim_{\varepsilon_{i_k}/\gamma}$, we have $\mu' \not\sim_{\varepsilon_{i_k}/\gamma} \nu_j$ for some $j \geq i_k$. This contradicts the fact that $\mu' \sim_{\varepsilon_j/\gamma} \nu_j$ and Lemma 6.1(3). Thus $\mu' R_{\varepsilon/\gamma} \nu'$ as required.

Finally, it is direct from definition that there exists a decreasing sequence $\{\varepsilon_i\}_i$ such that $\lim_i \varepsilon_i = D_b(\mu, \nu)$ and $\mu \sim_{\varepsilon_i} \nu$ for each i. Then the theorem follows. □

A direct consequence of the above theorem is that the bisimulation distance between two distributions vanishes if and only if they are bisimilar.

Corollary 6.1. *For any $\mu, \nu \in Dist(S)$, $\mu \sim \nu$ if and only if $D_b(\mu, \nu) = 0$.*

PROOF. Direct from Theorem 6.1, by noting that $\sim = \sim_0$. □

The next theorem shows that D_b is indeed a pseudo-metric.

Theorem 6.2. *The bisimulation distance D_b is a pseudo-metric on $Dist(S)$.*

PROOF. We need only to prove that D_b satisfies the triangle inequality

$$D_b(\mu, \nu) + D_b(\nu, \omega) \geq D_b(\mu, \omega).$$

By Theorem 6.1, we have $\mu \sim_{D_b(\mu,\nu)} \nu$ and $\nu \sim_{D_b(\nu,\omega)} \omega$. Then the result follows from Lemma 6.1(2). □

6.2. Modal Characterization of the Bisimulation Metric

We now present a Hennessy-Milner type modal logic motivated by [23, 25] to characterize the distance between distributions.

Definition 6.2. *The class $\mathcal{L}_m$ of modal formulae over AP, ranged over by φ, φ_1, φ_2, etc, is defined by the following grammar:*

$$\varphi \quad ::= \quad \mathcal{B} \mid \varphi \oplus p \mid \neg\varphi \mid \bigwedge_{i\in I} \varphi_i \mid \langle a \rangle \varphi$$

where $\mathcal{B} \subseteq 2^{AP}$, $p \in [0, 1]$, $a \in Act$, and I is an index set.

Given an input-enabled probabilistic automaton $\mathcal{A} = (S, Act, \to, L, \alpha)$ over AP, instead of defining the satisfaction relation $\models$ in the qualitative setting, the (discounted) semantics of the logic $\mathcal{L}_m$ is given in terms of functions from $Dist(S)$ to $[0,1]$. For any formula $\varphi \in \mathcal{L}_m$, the satisfaction function of φ, denoted by φ again for simplicity, is defined in a structurally inductive way as follows:

- $\mathcal{B}(\mu) := \sum_{A \in \mathcal{B}} \mu(A)$;
- $(\varphi \oplus p)(\mu) := \min\{\varphi(\mu) + p, 1\}$;
- $(\neg\varphi)(\mu) := 1 - \varphi(\mu)$;
- $(\bigwedge_{i\in I} \varphi_i)(\mu) := \inf_{i\in I} \varphi_i(\mu)$;
- $(\langle a\rangle\varphi)(\mu) := \sup_{\mu \xrightarrow{a} \mu'} \gamma \cdot \varphi(\mu')$.

Lemma 6.2. *For any $\varphi \in \mathcal{L}_m$, $\varphi : Dist(S) \to [0,1]$ is a continuous function.*

PROOF. We prove by induction on the structure of φ. The basis case when $\varphi \equiv \mathcal{B}$ is obvious. The case of $\varphi \equiv \varphi' \oplus p$, $\varphi \equiv \neg\varphi'$, and $\varphi \equiv \bigwedge_{i\in I} \varphi_i$ are all easy from induction. In the following we only consider the case when $\varphi \equiv \langle a\rangle\varphi'$.

Take arbitrarily $\{\mu_i\}_i$ with $\lim_i \mu_i = \mu$. We need to show there exists a subsequence $\{\mu_{i_k}\}_k$ of $\{\mu_i\}_i$ such that $\lim_k \varphi(\mu_{i_k}) = \varphi(\mu)$. Take arbitrarily $\varepsilon > 0$.

- Let $\mu^* \in Dist(S)$ such that $\mu \xrightarrow{a} \mu^*$ and $\varphi(\mu) \le \gamma \cdot \varphi'(\mu^*) + \varepsilon/2$. We have from the left-convergence of $\xrightarrow{a}$ that $\mu_i \xrightarrow{a} \nu_i$ for some ν_i, and $\lim_i \nu_i = \mu^*$. By induction, φ' is a continuous function. Thus we can find $N_1 \ge 1$ such that for any $i \ge N_1$, $|\varphi'(\mu^*) - \varphi'(\nu_i)| < \varepsilon/2\gamma$.
- For each $i \ge 1$, let $\mu_i^* \in Dist(S)$ such that $\mu_i \xrightarrow{a} \mu_i^*$ and $\varphi(\mu_i) \le \gamma \cdot \varphi'(\mu_i^*) + \varepsilon/2$. Then we have $\mu \xrightarrow{a} \nu^*$ with $\nu^* = \lim_k \mu_{i_k}^*$ for some convergent subsequence $\{\mu_{i_k}^*\}_k$ of $\{\mu_i^*\}_i$. Again, from the induction that φ' is continuous, we can find $N_2 \ge 1$ such that for any $k \ge N_2$, $|\varphi'(\mu_{i_k}^*) - \varphi'(\nu^*)| < \varepsilon/2\gamma$.

Let $N = \max\{N_1, N_2\}$. Then for any $k \ge N$, we have from $\mu \xrightarrow{a} \nu^*$ that

$$\begin{aligned} \varphi(\mu_{i_k}) - \varphi(\mu) &\le \gamma[\varphi'(\mu_{i_k}^*) - \varphi'(\nu^*)] + \gamma \cdot \varphi'(\nu^*) - \varphi(\mu) + \varepsilon/2 \\ &\le \gamma[\varphi'(\mu_{i_k}^*) - \varphi'(\nu^*)] + \varepsilon/2 < \varepsilon. \end{aligned}$$

Similarly, from $\mu_{i_k} \xrightarrow{a} \nu_{i_k}$ we have

$$\begin{aligned} \varphi(\mu) - \varphi(\mu_{i_k}) &\le \gamma[\varphi'(\mu^*) - \varphi'(\nu_{i_k})] + \gamma \cdot \varphi'(\nu_{i_k}) - \varphi(\mu_{i_k}) + \varepsilon/2 \\ &\le \gamma[\varphi'(\mu^*) - \varphi'(\nu_{i_k})] + \varepsilon/2 < \varepsilon. \end{aligned}$$

Thus $\lim_k \varphi(\mu_{i_k}) = \varphi(\mu)$ as required. □

From Lemma 6.2, and noting that the set $\{\mu' \mid \mu \xrightarrow{a} \mu'\}$ is compact for each μ and a, the supremum in the semantic definition of $\langle a\rangle\varphi$ can be replaced by maximum; that is, $(\langle a\rangle\varphi)(\mu) = \max_{\mu \xrightarrow{a} \mu'} \gamma \cdot \varphi(\mu')$. Now we define the logical distance for distributions.

Definition 6.3. *The* logic distance *of μ and ν in $Dist(S)$ of an input-enabled automaton is defined by*

$$D_l^{\mathcal{A}}(\mu, \nu) = \sup_{\varphi \in \mathcal{L}_m} |\varphi(\mu) - \varphi(\nu)| . \tag{2}$$

The logic distance for a general probabilistic automaton can be defined in terms of the input-enabled extension; that is, $D_l^{\mathcal{A}}(\mu,\nu) := D_l^{\mathcal{A}_\perp}(\mu,\nu)$. We always omit the superscripts for simplicity.

Now we can show that the logic distance exactly coincides with bisimulation distance for any pair of distributions.

Theorem 6.3. $D_b = D_l$.

PROOF. As both D_b and D_l are defined in terms of the input-enabled extension of automata, we only need to prove the result for input-enabled case. Let $\mu, \nu \in \mathit{Dist}(S)$. We first prove $D_b(\mu,\nu) \geq D_l(\mu,\nu)$. It suffices to show by structural induction that for any $\varphi \in \mathcal{L}_m$, $|\varphi(\mu) - \varphi(\nu)| \leq D_b(\mu,\nu)$. There are five cases to consider.

- $\varphi \equiv \mathcal{B}$ for some $\mathcal{B} \subseteq 2^{AP}$. Then $|\varphi(\mu) - \varphi(\nu)| = |\sum_{A\in\mathcal{B}}[\mu(A) - \nu(A)]| \leq d_{AP}(\mu,\nu) \leq D_b(\mu,\nu)$ by Theorem 6.1.
- $\varphi \equiv \varphi' \oplus p$. Assume without loss of generality that $\varphi'(\mu) \geq \varphi'(\nu)$. Then $\varphi(\mu) \geq \varphi(\nu)$. By induction, we have $\varphi'(\mu) - \varphi'(\nu) \leq D_b(\mu,\nu)$. Thus
$$|\varphi(\mu) - \varphi(\nu)| = \min\{\varphi'(\mu) + p, 1\} - \min\{\varphi'(\nu) + p, 1\} \leq \varphi'(\mu) - \varphi'(\nu) \leq D_b(\mu,\nu).$$
- $\varphi \equiv \neg\varphi'$. By induction, we have $|\varphi'(\mu) - \varphi'(\nu)| \leq D_b(\mu,\nu)$, thus $|\varphi(\mu) - \varphi(\nu)| = |1 - \varphi'(\mu) - 1 + \varphi'(\nu)| \leq D_b(\mu,\nu)$ as well.
- $\varphi \equiv \bigwedge_{i\in I}\varphi_i$. Assume $\varphi(\mu) \geq \varphi(\nu)$. For any $\varepsilon > 0$, let $j \in I$ such that $\varphi_j(\nu) \leq \varphi(\nu) + \varepsilon$. By induction, we have $|\varphi_j(\mu) - \varphi_j(\nu)| \leq D_b(\mu,\nu)$. Then
$$|\varphi(\mu) - \varphi(\nu)| \leq \varphi_j(\mu) - \varphi_j(\nu) + \varepsilon \leq D_b(\mu,\nu) + \varepsilon,$$
and $|\varphi(\mu) - \varphi(\nu)| \leq D_b(\mu,\nu)$ from the arbitrariness of ε.
- $\varphi \equiv \langle a\rangle\varphi'$. Assume $\varphi(\mu) \geq \varphi(\nu)$. Let $\mu'_* \in \mathit{Dist}(S)$ such that $\mu \xrightarrow{a} \mu'_*$ and $\gamma \cdot \varphi'(\mu'_*) = \varphi(\mu)$. From Theorem 6.1, we have $\mu \sim_{D_b(\mu,\nu)} \nu$. Thus there exists ν'_* such that $\nu \xrightarrow{a} \nu'_*$ and $\mu'_* \sim_{D_b(\mu,\nu)/\gamma} \nu'_*$. Hence $\gamma \cdot D_b(\mu'_*,\nu'_*) \leq D_b(\mu,\nu)$, and
$$|\varphi(\mu) - \varphi(\nu)| \leq \gamma \cdot [\varphi'(\mu'_*) - \varphi'(\nu'_*)] \leq \gamma \cdot D_b(\mu'_*,\nu'_*) \leq D_b(\mu,\nu)$$
where the second inequality is from induction.

Now we turn to the proof of $D_b(\mu,\nu) \leq D_l(\mu,\nu)$. We will achieve this by showing that the symmetric relations $R_\varepsilon = \{(\mu,\nu) \mid D_l(\mu,\nu) \leq \varepsilon\}$, where $\varepsilon \geq 0$, constitute an approximate bisimulation. Let $\mu R_\varepsilon \nu$ for some $\varepsilon \geq 0$. First, for any $\mathcal{B} \subseteq 2^{AP}$ we have
$$\left|\sum_{A\in\mathcal{B}}\mu(A) - \sum_{A\in\mathcal{B}}\nu(A)\right| = |\mathcal{B}(\mu) - \mathcal{B}(\nu)| \leq D_l(\mu,\nu) \leq \varepsilon.$$
Thus $d_{AP}(\mu,\nu) \leq \varepsilon$ as well. Now suppose $\mu \xrightarrow{a} \mu'$ for some μ'. We have to show that there is some ν' with $\nu \xrightarrow{a} \nu'$ and $D_l(\mu',\nu') \leq \varepsilon/\gamma$. Consider the set
$$\mathcal{K} = \{\omega \in \mathit{Dist}(S) \mid \nu \xrightarrow{a} \omega \text{ and } D_l(\mu',\omega) > \varepsilon/\gamma\}.$$
For each $\omega \in \mathcal{K}$, there must be some φ_ω such that $|\varphi_\omega(\mu') - \varphi_\omega(\omega)| > \varepsilon/\gamma$. As our logic includes the operator $\neg$, we can always assume that $\varphi_\omega(\mu') > \varphi_\omega(\omega) + \varepsilon/\gamma$. Let $p = \sup_{\omega\in\mathcal{K}}\varphi_\omega(\mu')$. Let
$$\varphi'_\omega = \varphi_\omega \oplus [p - \varphi_\omega(\mu')], \quad \varphi' = \bigwedge_{\omega\in\mathcal{K}}\varphi'_\omega, \quad \text{and} \quad \varphi = \langle a\rangle\varphi'.$$
Then from the assumption that $D_l(\mu,\nu) \leq \varepsilon$, we have $|\varphi(\mu) - \varphi(\nu)| \leq \varepsilon$. Furthermore, we check that for any $\omega \in \mathcal{K}$,
$$\varphi'_\omega(\mu') = \varphi_\omega(\mu') \oplus [p - \varphi_\omega(\mu')] = p.$$
Thus $\varphi(\mu) \geq \gamma \cdot \varphi'(\mu') = \gamma \cdot p$.

Let ν' be the distribution such that $\nu \xrightarrow{a} \nu'$ and $\varphi(\nu) = \gamma \cdot \varphi'(\nu')$. We are going to show that $\nu' \notin \mathcal{K}$, and then $D_l(\mu',\nu') \leq \varepsilon/\gamma$ as required. For this purpose, assume conversely that $\nu' \in \mathcal{K}$. Then
$$\begin{aligned}\varphi(\nu) &= \gamma \cdot \varphi'(\nu') \leq \gamma \cdot \varphi'_{\nu'}(\nu') \leq \gamma \cdot [\varphi_{\nu'}(\nu') + p - \varphi_{\nu'}(\mu')] \\ &< \gamma \cdot p - \varepsilon \leq \varphi(\mu) - \varepsilon,\end{aligned}$$
contradicting the fact that $|\varphi(\mu) - \varphi(\nu)| \leq \varepsilon$.

We have proven that $\{R_\varepsilon \mid \varepsilon \geq 0\}$ is an approximate bisimulation. Thus $\mu \sim_\varepsilon \nu$, and so $D_b(\mu,\nu) \leq \varepsilon$, whenever $D_l(\mu,\nu) \leq \varepsilon$. So we have $D_b(\mu,\nu) \leq D_l(\mu,\nu)$ from the arbitrariness of ε. □

6.3. A Fixed Point-Based Approach

In the following, we denote by $\mathcal{M}$ the set of pseudo-metrics over $Dist(S)$. Denote by $\mathbf{0}$ the zero pseudo-metric which assigns 0 to each pair of distributions. For any $d, d' \in \mathcal{M}$, we write $d \leq d'$ if $d(\mu,\nu) \leq d'(\mu,\nu)$ for any μ and ν. Obviously $\leq$ is a partial order, and $(\mathcal{M}, \leq)$ is a complete lattice.

Definition 6.4. *Let $\mathcal{A} = (S, Act, \to, L, \alpha)$ be an input-enabled probabilistic automaton. We define the function $F : \mathcal{M} \to \mathcal{M}$ as follows. For any $\mu, \nu \in Dist(S)$,*

$$F(d)(\mu,\nu) = \max_{a \in Act} \quad \{ \quad d_{AP}(\mu,\nu),\ \sup_{\mu \xrightarrow{a} \mu'} \inf_{\nu \xrightarrow{a} \nu'} \gamma \cdot d(\mu',\nu'),\ \sup_{\nu \xrightarrow{a} \nu'} \inf_{\mu \xrightarrow{a} \mu'} \gamma \cdot d(\mu',\nu') \}.$$

Then, F is monotonic with respect to $\leq$, and by Knaster-Tarski theorem, F has a least fixed point, denoted $D_f^{\mathcal{A}}$, given by

$$D_f^{\mathcal{A}} = \bigvee_{n=0}^{\infty} F^n(\mathbf{0}) .$$

Here $\bigvee$ means the supremum with respect to the order $\leq$.

Once again, the fixed point-based distance for a general probabilistic automaton can be defined in terms of the input-enabled extension; that is, $D_f^{\mathcal{A}}(\mu,\nu) := D_f^{\mathcal{A}_\perp}(\mu,\nu)$. We always omit the superscripts for simplicity.

Similar to Lemma 6.2, we can show that the supremum (resp. infimum) in Definition 6.4 can be replaced by maximum (resp. minimum). Now we show that D_f coincides with D_b.

Theorem 6.4. $D_f = D_b$.

As both D_f and D_b are defined in terms of the input-enabled extension of automata, we only need to prove Theorem 6.4 for input-enabled case, which will be obtained by combining Lemma 6.3 and Lemma 6.6 below.

Lemma 6.3. *For input-enabled probabilistic automata, $D_f \leq D_b$.*

PROOF. It suffices to prove by induction that for any $n \geq 0$, $F^n(\mathbf{0}) \leq D_b$. The case of $n = 0$ is trivial. Suppose $F^n(\mathbf{0}) \leq D_b$ for some $n \geq 0$. Then for any $a \in Act$ and any μ, ν, we have

(1) $d_{AP}(\mu,\nu) \leq D_b(\mu,\nu)$ by the fact that $\mu \sim_{D_b(\mu,\nu)} \nu$;

(2) Note that $\mu \sim_{D_b(\mu,\nu)} \nu$. Whenever $\mu \xrightarrow{a} \mu'$, we have $\nu \xrightarrow{a} \nu'$ for some ν' such that $\mu' \sim_{D_b(\mu,\nu)/\gamma} \nu'$, and hence $\gamma \cdot D_b(\mu',\nu') \leq D_b(\mu,\nu)$. That, together with the assumption $F^n(\mathbf{0}) \leq D_b$, implies

$$\max_{\mu \xrightarrow{a} \mu'} \min_{\nu \xrightarrow{a} \nu'} \gamma \cdot F^n(\mathbf{0})(\mu',\nu') \leq D_b(\mu,\nu).$$

The symmetric form can be similarly proved.

Summing up (1) and (2), we have $F^{n+1}(\mathbf{0}) \leq D_b$. □

To prove the other direction, we first introduce the notion of bounded approximate bisimulations.

Definition 6.5. *Let $\mathcal{A}$ be an input-enabled probabilistic automaton. We define symmetric relations*

- $\overset{\varepsilon}{\sim}_0 := Dist(S) \times Dist(S)$ *for any* $\varepsilon \geq 0$;
- *for $n \geq 0$, $\mu \overset{\varepsilon}{\sim}_{n+1} \nu$ if $d_{AP}(\mu,\nu) \leq \varepsilon$ and whenever $\mu \xrightarrow{a} \mu'$, there exists $\nu \xrightarrow{a} \nu'$ for some ν' such that $\mu' \overset{\varepsilon/\gamma}{\sim}_n \nu'$.*
- $\overset{\varepsilon}{\sim} := \bigcap_{n \geq 0} \overset{\varepsilon}{\sim}_n$.

The following lemma collects some useful properties of $\overset{\varepsilon}{\sim}_n$ and $\overset{\varepsilon}{\sim}$.

Lemma 6.4. *1.* $\overset{\varepsilon}{\sim}_n \subseteq \overset{\varepsilon}{\sim}_m$ *provided that* $n \geq m$*;*

2. for any $n \geq 0$, $\overset{\varepsilon}{\sim}_n \subseteq \overset{\varepsilon'}{\sim}_n$ *provided that* $\varepsilon \leq \varepsilon'$*;*

3. for any $n \geq 0$, $\overset{\varepsilon}{\sim}_n$ *is continuous;*

4. $\overset{\varepsilon}{\sim} = \sim_\varepsilon$.

PROOF. Items 1, 2, and 3 are easy by induction, and so is the $\sim_\varepsilon \subseteq \overset{\varepsilon}{\sim}$ part of item 4. To prove $\overset{\varepsilon}{\sim} \subseteq \sim_\varepsilon$, we show that $\{\overset{\varepsilon}{\sim} | \ \varepsilon \geq 0\}$ is an approximate bisimulation. Suppose $\mu \overset{\varepsilon}{\sim} \nu$. Then $d_{AP}(\mu, \nu) \leq \varepsilon$ by definition. Now let $\mu \xrightarrow{a} \mu'$. For each $n \geq 0$, from the assumption that $\mu \overset{\varepsilon}{\sim}_{n+1} \nu$ we have $\nu \xrightarrow{a} \nu_n$ such that $\mu' \overset{\varepsilon/\gamma}{\sim}_n \nu_n$. Let $\{\nu_{i_k}\}_k$ be a convergent subsequence of $\{\nu_n\}_n$ such that $\lim_k \nu_{i_k} = \nu'$ for some ν'. Then from the continuity of $\xrightarrow{a}$ we have $\nu \xrightarrow{a} \nu'$. We claim further that $\mu' \overset{\varepsilon/\gamma}{\sim} \nu'$. Otherwise there exists N such that $\mu' \not\sim \varepsilon/\gamma_N \nu'$. Now by the continuity of $\overset{\varepsilon/\gamma}{\sim}_N$, we have $\mu' \not\sim \varepsilon/\gamma_N \nu_j$ for some $j \geq N$. This contradicts the fact that $\mu' \overset{\varepsilon/\gamma}{\sim}_j \nu_j$ and item 1. □

Lemma 6.5. *For any* $n \geq 0$*, we have* $\mu \overset{F^n(\mathbf{0})(\mu,\nu)}{\sim}_n \nu$.

PROOF. We prove this lemma by induction on n. The case of $n = 0$ is trivial. Suppose $\mu \overset{F^n(\mathbf{0})(\mu,\nu)}{\sim}_n \nu$ for some $n \geq 0$. Let $a \in Act$. By definition, we have

$$F^{n+1}(\mathbf{0})(\mu, \nu) \geq \max_{\mu \xrightarrow{a} \mu'} \min_{\nu \xrightarrow{a} \nu'} \gamma \cdot F^n(\mathbf{0})(\mu', \nu').$$

Thus for any $\mu \xrightarrow{a} \mu'$, there exists $\nu \xrightarrow{a} \nu'$ such that $\gamma \cdot F^n(\mathbf{0})(\mu', \nu') \leq F^{n+1}(\mathbf{0})(\mu, \nu)$. By induction, we know $\mu' \overset{F^n(\mathbf{0})(\mu',\nu')}{\sim}_n \nu'$, thus $\mu' \overset{F^{n+1}(\mathbf{0})(\mu,\nu)/\gamma}{\sim}_n \nu'$ from Lemma 6.4(2). On the other hand, we have $F^{n+1}(\mathbf{0})(\mu, \nu) \geq d_{AP}(\mu, \nu)$ by definition. Thus we have $\mu \overset{F^{n+1}(\mathbf{0})(\mu,\nu)}{\sim}_{n+1} \nu$. □

With the two lemmas above, we can prove that $D_b \leq D_f$.

Lemma 6.6. *For input-enabled probabilistic automata,* $D_b \leq D_f$.

PROOF. For any μ and ν, by Lemmas 6.5 and 6.4(2), we have $\mu \overset{D_f(\mu,\nu)}{\sim}_n \nu$ for all $n \geq 0$, so $\mu \overset{D_f(\mu,\nu)}{\sim} \nu$ by definition. Then from Lemma 6.4(4) we have $\mu \sim_{D_f(\mu,\nu)} \nu$, hence $D_b(\mu, \nu) \leq D_f(\mu, \nu)$. □

6.4. Comparison with State-Based Metric

In this subsection, we show that our distribution-based bisimulation metric is upper bounded by the state-based game bisimulation metric [24] for MDPs. This game bisimulation metric is particularly attractive as it preserves probabilistic reachability, long-run, and discounted average behaviours [50]. We first recall the definition of state-based game bisimulation metric for MDPs in [24]:

Definition 6.6. *Given* $\mu, \nu \in Dist(S)$, $\mu \otimes \nu$ *is defined as the set of* weight *functions* $\lambda : S \times S \to [0, 1]$ *such that for any* $s, t \in S$,

$$\sum_{s \in S} \lambda(s, t) = \nu(t) \quad \textit{and} \quad \sum_{t \in S} \lambda(s, t) = \mu(s).$$

Given a metric d *defined on* S, *we lift it to* $Dist(S)$ *by defining*

$$d(\mu, \nu) = \inf_{\lambda \in \mu \otimes \nu} \left(\sum_{s,t \in S} \lambda(s, t) \cdot d(s, t) \right).$$

Actually the infimum in the above definition is attainable.

Definition 6.7. *We define the function $f : \mathcal{M} \to \mathcal{M}$ as follows. For any $s, t \in S$,*

$$f(d)(s,t) = \max_{a \in Act} \left\{ 1 - \delta_{L(s),L(t)},\ \sup_{s \xrightarrow{a}_{\mathsf{P}} \mu} \inf_{t \xrightarrow{a}_{\mathsf{P}} \nu} \gamma \cdot d(\mu,\nu),\ \sup_{t \xrightarrow{a}_{\mathsf{P}} \nu} \inf_{s \xrightarrow{a}_{\mathsf{P}} \mu} \gamma \cdot d(\mu,\nu) \right\}$$

where $\delta_{L(s),L(t)} = 1$ if $L(s) = L(t)$, and 0 otherwise. We take $\inf \emptyset = 1$ and $\sup \emptyset = 0$. Again, f is monotonic with respect to $\leq$, and by Knaster-Tarski theorem, F has a least fixed point, denoted d_f, given by

$$d_f = \bigvee_{n=0}^{\infty} f^n(\mathbf{0}).$$

Now we can prove the quantitative extension of Lemma 5.1. Without loss of any generality, we assume that $\mathcal{A}$ itself is input-enabled. Let $d_n = f^n(\mathbf{0})$ and $D_n = F^n(\mathbf{0})$ in Definition 6.4.

Lemma 6.7. *For any $n \geq 1$, $d_{AP}(\mu,\nu) \leq d_n(\mu,\nu)$.*

PROOF. Let λ be the weight function such that $d_n(\mu,\nu) = \sum_{s,t \in S} \lambda(s,t) \cdot d_n(s,t)$. Since $d_n(s,t) \geq 1 - \delta_{L(s),L(t)}$, we have

$$d_n(\mu,\nu) \geq 1 - \sum_{s,t:L(s)=L(t)} \lambda(s,t).$$

On the other hand, for any $A \subseteq AP$, recall that $S(A) = \{s \in S \mid L(s) = A\}$. Then

$$\begin{aligned} \mu(A) - \nu(A) &= \sum_{s \in S(A)} \mu(s) - \sum_{t \in S(A)} \nu(t) \\ &= \sum_{s \in S(A)} \sum_{t \notin S(A)} \lambda(s,t) - \sum_{t \in S(A)} \sum_{s \notin S(A)} \lambda(s,t). \end{aligned}$$

Let $\mathcal{B} \subseteq 2^{AP}$ such that $d_{AP}(\mu,\nu) = \sum_{A \in \mathcal{B}} [\mu(A) - \nu(A)]$. Then

$$d_{AP}(\mu,\nu) \leq \sum_{A \in \mathcal{B}} \sum_{s \in S(A)} \sum_{t \notin S(A)} \lambda(s,t) \leq \sum_{s,t:L(s) \neq L(t)} \lambda(s,t),$$

and the result follows. □

Theorem 6.5. *Let $\mathcal{A}$ be a probabilistic automaton. Then $D_f \leq d_f$.*

PROOF. We prove by induction on n that $D_n(\mu,\nu) \leq d_n(\mu,\nu)$ for any $\mu, \nu \in Dist(S)$ and $n \geq 0$. The case $n = 0$ is obvious. Suppose the result holds for some $n - 1 \geq 0$. Then from Lemma 6.7, we need only to show that for any $\mu \xrightarrow{a} \mu'$ there exists $\nu \xrightarrow{a} \nu'$ such that $\gamma \cdot D_{n-1}(\mu',\nu') \leq d_n(\mu,\nu)$.

Let $\mu \xrightarrow{a} \mu'$. Then for any $s \in S$, $s \xrightarrow{a}_{\mathsf{P}} \mu_s$ with $\mu' = \sum_{s \in S} \mu(s) \cdot \mu_s$. By definition of d_n, for any $t \in S$, we have $t \xrightarrow{a}_{\mathsf{P}} \nu_t$ such that $\gamma \cdot d_{n-1}(\mu_s,\nu_t) \leq d_n(s,t)$. Thus $\nu \xrightarrow{a} \nu' := \sum_{t \in S} \nu(t) \cdot \nu_t$, and by induction, $D_{n-1}(\mu',\nu') \leq d_{n-1}(\mu',\nu')$. Now it suffices to prove $\gamma \cdot d_{n-1}(\mu',\nu') \leq d_n(\mu,\nu)$.

Let $\lambda \in \mu \otimes \nu$ and $\gamma_{s,t} \in \mu_s \otimes \nu_t$ be the weight functions such that

$$d_n(\mu,\nu) = \sum_{s,t \in S} \lambda(s,t) \cdot d_n(s,t),\ \ d_{n-1}(\mu_s,\nu_t) = \sum_{u,v \in S} \gamma_{s,t}(u,v) \cdot d_{n-1}(u,v).$$

Then

$$\begin{aligned} d_n(\mu,\nu) &\geq \gamma \cdot \sum_{s,t \in S} \lambda(s,t) \cdot d_{n-1}(\mu_s,\nu_t) \\ &= \gamma \cdot \sum_{u,v \in S} \sum_{s,t \in S} \lambda(s,t) \gamma_{s,t}(u,v) \cdot d_{n-1}(u,v). \end{aligned}$$

We need to show that the function $\eta(u,v) := \sum_{s,t\in S} \lambda(s,t)\gamma_{s,t}(u,v)$ is a weight function for μ' and ν'. Indeed, it is easy to check that

$$\begin{aligned}\sum_u \eta(u,v) &= \sum_{s,t\in S} \lambda(s,t)\sum_u \gamma_{s,t}(u,v) = \sum_{s,t\in S} \lambda(s,t)\nu_t(v)\\ &= \sum_{t\in S} \nu(t)\nu_t(v) = \nu'(v).\end{aligned}$$

Similarly, we have $\sum_v \eta(u,v) = \mu'(u)$. □

Example 6.1. *Consider Fig. 1, and assume $\varepsilon_1 \geq \varepsilon_2 \geq 0$ and $\gamma = 1$. It is easy to check that $D_f(\delta_q, \delta_{q'}) = \frac{1}{2}(\varepsilon_1 - \varepsilon_2)$. However, according to the state-based bisimulation metric, $d_f(\delta_{r_1}, \delta_{r'}) = \frac{1}{6} + \varepsilon_1$ and $d_f(\delta_{r_2}, \delta_{r'}) = \frac{1}{6} + \varepsilon_2$. Thus $d_f(\delta_q, \delta_{q'}) = \frac{1}{6} + \frac{1}{2}(\varepsilon_1 + \varepsilon_2)$.*

6.5. Comparison with Equivalence Metric

Note that we can easily extend the equivalence relation defined in Definition 4.5 to a notion of equivalence metric:

Definition 6.8 (Equivalence Metric). *Let $\mathcal{A}_1$ and $\mathcal{A}_2$ be two reactive automata with the same set of actions Act. We say $\mathcal{A}_1$ and $\mathcal{A}_2$ are ε-equivalent, denoted $\mathcal{A}_1 \sim_\varepsilon^d \mathcal{A}_2$, if for any input word $w = a_1a_2\ldots a_n \in Act^*$, $|\mathcal{A}_1(w) - \mathcal{A}_2(w)| \leq \varepsilon$. Furthermore, the equivalence distance between $\mathcal{A}_1$ and $\mathcal{A}_2$ is defined by $D_d(\mathcal{A}_1, \mathcal{A}_2) := \inf\{\varepsilon \geq 0 \mid \mathcal{A}_1 \sim_\varepsilon^d \mathcal{A}_2\}$.*

Now we show that for reactive automata, the equivalence metric D_d coincide with our un-discounted bisimulation metric D_b, which may be regarded as a quantitative extension of Lemma 5.2.

Proposition 6.1. *Let $\mathcal{A}_1$ and $\mathcal{A}_2$ be two reactive automata with the same set of actions Act. Let the discounting factor $\gamma = 1$. Then $D_d(\mathcal{A}_1, \mathcal{A}_2) = D_b(\alpha_1, \alpha_2)$ where D_b is defined in the direct sum of $\mathcal{A}_1$ and $\mathcal{A}_2$.*

PROOF. We first show that $D_d(\mathcal{A}_1, \mathcal{A}_2) \leq D_b(\alpha_1, \alpha_2)$. For each input word $w = a_1a_2\ldots a_n$, it is easy to check that $\mathcal{A}_i(w) = \varphi(\alpha_i)$ where $\varphi = \langle a_1\rangle\langle a_2\rangle\ldots\langle a_n\rangle(F_1 \cup F_2)$. As we have shown that $D_b = D_l$, it holds $|\mathcal{A}_1(w) - \mathcal{A}_2(w)| \leq D_b(\alpha_1, \alpha_2)$, and hence $\mathcal{A}_1 \sim_{D_b(\alpha_1,\alpha_2)}^d \mathcal{A}_2$. Then $D_d(\mathcal{A}_1, \mathcal{A}_2) \leq D_b(\alpha_1, \alpha_2)$ by definition.

Now we turn to the proof of $D_d(\mathcal{A}_1, \mathcal{A}_2) \geq D_b(\alpha_1, \alpha_2)$. First we show that

$$R_\varepsilon = \{(\mu,\nu) \mid \mu \in Dist(S_1), \nu \in Dist(S_2), \mathcal{A}_1^\mu \sim_\varepsilon^d \mathcal{A}_2^\nu\}$$

is an approximate bisimulation. Here for a probabilistic automaton $\mathcal{A}$, we denote by $\mathcal{A}^\mu$ the automaton which is the same as $\mathcal{A}$ except that the initial distribution is replaced by μ. Let $\mu R_\varepsilon \nu$. Since $L(s) \in \{\emptyset, AP\}$ for all $s \in S_1 \cup S_2$, we have $\mu(AP) + \mu(\emptyset) = \nu(AP) + \nu(\emptyset) = 1$. Thus

$$d_{AP}(\mu,\nu) = |\mu(AP) - \nu(AP)| = |\mu(F_1) - \nu(F_2)|.$$

Note that $\mu(F_1) = \mathcal{A}_1^\mu(e)$ and $\nu(F_2) = \mathcal{A}_2^\nu(e)$, where e is the empty string. Then $d_{AP}(\mu,\nu) = |\mathcal{A}_1^\mu(e) - \mathcal{A}_2^\nu(e)| \leq \varepsilon$.

Let $\mu \xrightarrow{a} \mu'$ and $\nu \xrightarrow{a} \nu'$. We need to show $\mu' R_\varepsilon \nu'$, that is, $\mathcal{A}_1^{\mu'} \sim_\varepsilon^d \mathcal{A}_2^{\nu'}$. For any $w \in Act^*$ and $i = 1, 2$, note that $\mathcal{A}_i^{\mu'}(w) = \mathcal{A}_i^\mu(aw)$. Then

$$|\mathcal{A}_1^{\mu'}(w) - \mathcal{A}_2^{\nu'}(w)| = |\mathcal{A}_1^\mu(aw) - \mathcal{A}_2^\nu(aw)| \leq \varepsilon,$$

and hence $\mathcal{A}_1^{\mu'} \sim_\varepsilon^d \mathcal{A}_2^{\nu'}$ as required.

Having proven that R_ε is an approximate bisimulation, we know $\mathcal{A}_1 \sim_\varepsilon^d \mathcal{A}_2$ implies $\alpha_1 \sim_\varepsilon \alpha_2$. Thus

$$D_d(\mathcal{A}_1, \mathcal{A}_2) = \inf\{\varepsilon \mid \mathcal{A}_1 \sim_\varepsilon^d \mathcal{A}_2\} \geq \inf\{\varepsilon \mid \alpha_1 \sim_\varepsilon \alpha_2\} = D_b(\alpha_1, \alpha_2).$$

□

7. Comparison with Distribution-based Bisimulations in Literature

In this section, we review some distribution-based definitions of bisimulation in the literature and discuss their relations. We first recall the definition in [20] except that we focus on its strong counterpart. For this, we need some notations. Recall that $EA(s)$ denotes the set of actions which can be performed in s. A distribution μ is *consistent*, denoted $\overrightarrow{\mu}$, if $EA(s) = EA(t)$ for any $s, t \in supp(\mu)$, i.e., all states in the support of μ have the same set of enabled actions. In case μ is consistent, we also let $EA(\mu) = EA(s)$ for some $s \in supp(\mu)$.

Definition 7.1. *Let $\mathcal{A} = (S, Act, \rightarrow, L, \alpha)$ be a probabilistic automaton. A symmetric relation $R \subseteq Dist(S) \times Dist(S)$ is a §-bisimulation if $\mu R \nu$ implies that*

1. *$\mu(A) = \nu(A)$ for each $A \subseteq AP$,*
2. *for each $a \in Act$ whenever $\mu \xrightarrow{a} \mu'$, there exists a transition $\nu \xrightarrow{a} \nu'$ such that $\mu' R \nu'$, and*
3. *if not $\overrightarrow{\mu}$, there exist convex decompositions $\mu = \sum_{1 \le i \le n} p_i \cdot \mu_i$ and $\nu = \sum_{1 \le i \le n} p_i \cdot \nu_i$ such that $\overrightarrow{\mu_i}$ and $\mu_i R \nu_i$ for each $1 \le i \le n$.*

We write $\mu \sim_{\S}^{\mathcal{A}} \nu$ if there is a §-bisimulation R such that $\mu R \nu$.

Note that Definition 7.1 is given directly for general probabilistic automata without the need for input-enabled extension. Recently, another definition of bisimulation based on distributions was introduced in [21]. The main difference arises in the lifting transition relation of distributions. Let $\mathbb{A} \subseteq Act$ and $S_{\mathbb{A}} = \{s \in S \mid EA(s) \cap \mathbb{A} \neq \emptyset\}$, i.e., $S_{\mathbb{A}}$ contains all states which is able to perform an action in $\mathbb{A}$. Instead of labelling a transition with a single action, transitions of states and distributions in [21] are labelled by a set of actions, denoted $\overset{\mathbb{A}}{\rightsquigarrow}$. Formally, $s \overset{\mathbb{A}}{\rightsquigarrow} \mu$ if there exists $a \in \mathbb{A}$ such that $s \overset{a}{\rightsquigarrow} \mu$, otherwise we say actions $\mathbb{A}$ are blocked at s. Accordingly, $\mu \overset{\mathbb{A}}{\rightsquigarrow} \nu$ if $\mu(S_{\mathbb{A}}) > 0$ and for each $s \in S_{\mathbb{A}} \cap supp(\mu)$, there exists $s \overset{\mathbb{A}}{\rightsquigarrow} \mu_s$ such that

$$\nu = \frac{1}{\mu(S_{\mathbb{A}})} \sum_{s \in S_{\mathbb{A}} \cap supp(\mu)} \mu(s) \cdot \mu_s.$$

Intuitively, a distribution μ is able to perform a transition with label $\mathbb{A}$ if and only if at least one of its supports can perform an action in $\mathbb{A}$. Furthermore, all states in $S_{\mathbb{A}} \cap supp(\mu)$ should perform such a transition in the meanwhile. The resulting distribution is the weighted sum of all the resulting distributions with weights equal to their probabilities in μ. Since it may happen that some states in $supp(\mu)$ cannot perform such a transition, i.e., $supp(\mu) \nsubseteq S_{\mathbb{A}}$, we need the normalizer $\frac{1}{\mu(S_{\mathbb{A}})}$ in order to obtain a valid distribution. Below follows the definition of bisimulation in [21], where $\mu(\mathbb{A}, A) = \mu(\{s \in S_{\mathbb{A}} \mid L(s) = A\})$, the probability of states in μ labelled by A while being able to perform actions in $\mathbb{A}$.

Definition 7.2. *Let $\mathcal{A} = (S, Act, \rightarrow, L, \alpha)$ be a probabilistic automaton. A symmetric relation $R \subseteq Dist(S) \times Dist(S)$ is a †-bisimulation if $\mu R \nu$ implies that*

1. *$\mu(\mathbb{A}, A) = \nu(\mathbb{A}, A)$ for each $\mathbb{A} \subseteq Act$ and $A \subseteq AP$,*
2. *for each $\mathbb{A} \subseteq Act$, whenever $\mu \overset{\mathbb{A}}{\rightsquigarrow} \mu'$, there exists a transition $\nu \overset{\mathbb{A}}{\rightsquigarrow} \nu'$ such that $\mu' R \nu'$.*

We write $\mu \sim_{\dagger}^{\mathcal{A}} \nu$ if there is a †-bisimulation R such that $\mu R \nu$.

For simplicity, we omit the superscript $\mathcal{A}$ of all relations if it is clear from the context. Similar to Theorem 5.1, we can prove that both $\sim_{\S}$ and $\sim_{\dagger}$ are linear. In the following, we show that in general $\sim$ is strictly coarser than $\sim_{\S}$ and $\sim_{\dagger}$, while $\sim_{\S}$ and $\sim_{\dagger}$ are incomparable. Furthermore, if restricted to input-enabled probabilistic automata, $\sim$ and $\sim_{\S}$ coincide.

Theorem 7.1. *Let $\mathcal{A} = (S, Act, \rightarrow, L, \alpha)$ be a probabilistic automaton.*

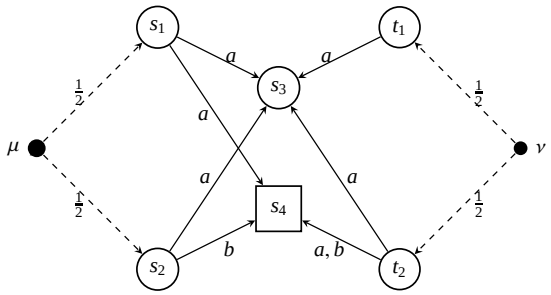

Figure 2: $\mu \sim \nu$ but $\mu \nsim_{\S} \nu$ and $\mu \nsim_{\dagger} \nu$.

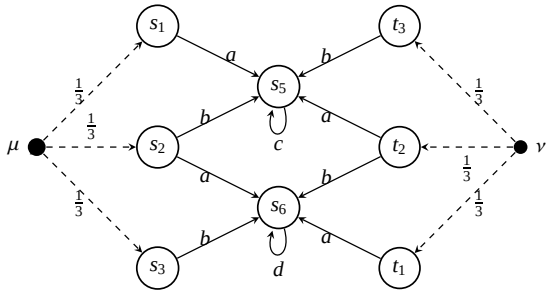

Figure 3: $\mu \sim_{\dagger} \nu$ but $\mu \nsim_{\S} \nu$.

1. $\sim_{\dagger}$, $\sim_{\S} \subsetneq \sim$.
2. $\sim_{\dagger}$ *and* $\sim_{\S}$ *are incomparable.*
3. *If* $\mathcal{A}$ *is input-enabled,* $\sim_{\dagger} \subsetneq \sim_{\S} \equiv \sim$.

PROOF. We prove the theorem in several steps:

1. $\sim_{\S} \subseteq \sim$. Let
$$R = \{(p \cdot \mu_1 + (1-p) \cdot \delta_{\perp}, p \cdot \nu_1 + (1-p) \cdot \delta_{\perp}) \mid p \in [0,1] \wedge \mu_1 \sim_{\S}^{\mathcal{A}} \nu_1\}.$$
It suffices to show that R is a bisimulation in $\mathcal{A}_{\perp}$. Let $\mu R \nu$ such that $\mu \equiv (p \cdot \mu_1 + (1-p) \cdot \delta_{\perp})$ and $\nu \equiv (p \cdot \nu_1 + (1-p) \cdot \delta_{\perp})$ with $\mu_1 \sim_{\S}^{\mathcal{A}} \nu_1$. It is easy to show that $\mu(A) = \nu(A)$ for any $A \subseteq AP$. Suppose $\mu \xrightarrow{a}_{\perp} \mu'$ for some $a \in Act$, we shall show that there exists $\nu \xrightarrow{a}_{\perp} \nu'$ such that $\mu' R \nu'$. This is trivial if $a \notin EA(s)$ for all $s \in supp(\mu_1)$, or $a \in EA(\mu_1)$ and $\overrightarrow{\mu_1}$. Now let $\mu_1 \equiv \sum_{i \in I} p_i \cdot \mu_i$ with μ_i being consistent for each $i \in I$. Let $J = \{i \in I \mid a \in EA(\mu_i)\}$. By Definition 5.3, for each $i \in J$ there exists μ_i' such that $\mu_i \xrightarrow{a} \mu_i'$ and
$$\mu' \equiv \sum_{i \in J} p_i \cdot \mu_i' + (1 - \sum_{i \in J} p_i) \cdot \delta_{\perp}.$$
Since $\mu_1 \sim_{\S}^{\mathcal{A}} \nu_1$, there exists $\nu_1 \equiv \sum_{i \in I} p_i \cdot \nu_i$ such that $\overrightarrow{\nu_i}$ and $\mu_i \sim_{\S}^{\mathcal{A}} \nu_i$ for each $i \in I$. Moreover, for every $i \in J$, there exists $\nu_i \xrightarrow{a} \nu_i'$ such that $\mu_i' \sim_{\S}^{\mathcal{A}} \nu_i'$. Now let
$$\nu' \equiv \sum_{i \in J} p_i \cdot \nu_i' + (1 - \sum_{i \in J} p_i) \cdot \delta_{\perp}.$$
Then $\nu \xrightarrow{a}_{\perp} \nu'$. From the linearity of $\sim_{\S}^{\mathcal{A}}$ and the definition of R, we can easily show that $\mu' R \nu'$, thus R is a bisimulation in $\mathcal{A}_{\perp}$.

2. $\sim_\dagger \subseteq \sim$. The proof is similar as the above case. Let

$$R = \{(p \cdot \mu_1 + (1-p) \cdot \delta_\perp, p \cdot \nu_1 + (1-p) \cdot \delta_\perp) \mid p \in [0,1] \wedge \mu_1 \sim_\dagger^{\mathcal{A}} \nu_1\}.$$

Then we show that R is a bisimulation in $\mathcal{A}_\perp$. Let $\mu R \nu$, i.e., there exists $p \in [0,1]$, μ_1, and ν_1 such that $\mu \equiv p \cdot \mu_1 + (1-p) \cdot \delta_\perp$, $\nu \equiv p \cdot \nu_1 + (1-p) \cdot \delta_\perp$, and $\mu_1 \sim_\dagger^{\mathcal{A}} \nu_1$. Apparently, $\mu(A) = \nu(A)$ for any $A \subseteq AP$. Let $\mu \xrightarrow{a}_\perp \mu'$. We shall show there exists $\nu \xrightarrow{a}_\perp \nu'$ such that $\mu' R \nu'$. Note $\mu \xrightarrow{a}_\perp \mu'$ indicates that $\mu_1 \overset{\{a\}}{\rightsquigarrow} \mu_1'$ with $\mu' \equiv p \cdot \mu_1' + (1-p) \cdot \delta_\perp$. Since $\mu_1 \sim_\dagger^{\mathcal{A}} \nu_1$, there exists $\nu_1 \overset{\{a\}}{\rightsquigarrow} \nu_1'$ such that $\mu_1' \sim_\dagger^{\mathcal{A}} \nu_1'$. Therefore there exists $\nu \xrightarrow{a}_\perp \nu'$ with $\nu' \equiv p \cdot \nu_1' + (1-p) \cdot \delta_\perp$. By the definition of R, $\mu' R \nu'$ as desired.

3. $\sim \not\subseteq \sim_\S$ and $\sim \not\subseteq \sim_\dagger$. Let μ and ν be two distributions as in Fig. 2, where each state is labelled by its shape. By adding extra transitions to the dead state, we can see $\mu \sim \nu$ by showing that the following relation is a bisimulation: $\{(\mu, \nu), (\nu, \mu)\} \cup ID$, where ID denotes the identity relation. However, neither $\mu \sim_\S \nu$ nor $\mu \sim_\dagger \nu$ holds. For the former, since μ is not consistent, we shall split it to δ_{s_1} and δ_{s_2} by Definition 7.1, where δ_{s_1} cannot be simulated by ν and its successors. To see $\mu \nsim_\dagger \nu$, let $\mathbb{A} = \{a, b\}$. Then μ can evolve into box states with probability 1, while the probability is at most 0.5 for ν.

4. $\sim_\dagger \not\subseteq \sim_\S$. Let μ and ν be two distributions as in Fig. 3, where all states have the same label. Let $R = \{(\mu, \nu), (\nu, \mu)\} \cup ID$. By Definition 7.2, it is easy to see that R is a †-bisimulation. Therefore $\mu \sim_\dagger \nu$. However, $\mu \nsim_\S \nu$. Since s_1, s_2, and s_3 have different enabled actions, thus μ shall be split into three dirac distributions, none of which can be simulated by any successor of ν.

5. $\sim_\S \not\subseteq \sim_\dagger$. Let μ and ν be the distributions in Fig. 2 except that s_1 and t_1 have a transition with label b to some state with a different label from all the others. We see that $\mu \sim_\S \nu$, since by adding transitions with label b to s_1 and t_1, both μ and ν are consistent, thus need not to be split. However, $\mu \nsim_\dagger \nu$. To see this, let $\mathbb{A} = \{a, b\}$. Then μ can evolve into box states via a transition with label $\mathbb{A}$ with probability 1, which is not possible in ν.

6. If $\mathcal{A}$ is input-enabled, then $\sim \subseteq \sim_\S$. Since in an input-enabled probabilistic automaton, all distributions are consistent, and there is no need to split, i.e., the last condition in Definition 7.1 is redundant. The counterexample given in the above case can be used to show that $\sim_\S$ is strictly coarser than $\sim_\dagger$ even if restricted to input-enabled probabilistic automata. □

7.1. Bisimulations and Trace Equivalences

In this subsection, we discuss how different bisimulation relations and trace equivalences are related in Segala's automata. A path $\sigma \in S \times (Act \times S)^*$ is an alternative sequence of states and actions, and a trace $w \in Act^*$ is a sequence of actions. Let $Paths^*(\mathcal{A})$ denote the set of all finite paths of a probabilistic automaton $\mathcal{A}$ and $\sigma \downarrow$ the last state of σ. Due to the non-determinism in a probabilistic automaton, a *scheduler* is often adopted in order to obtain a fully probabilistic system. A scheduler can be seen as a function taking a history execution as input, while choosing a transition as the next step for the current state. Formally,

Definition 7.3. *Let $\mathcal{A} = (S, Act, \rightarrow, L, \alpha)$ be a probabilistic automaton. A scheduler $\pi : Paths^*(\mathcal{A}) \mapsto Dist(Act \times Dist(S))$ of $\mathcal{A}$ is a function such that $\pi(\sigma)(a, \mu) > 0$ only if $(\sigma\downarrow, a, \mu) \in \rightarrow$.*

Let $\mathcal{A} = (S, Act, \rightarrow, L, \alpha)$ be a probabilistic automaton, π a scheduler, w a trace of $\mathcal{A}$, and μ a distribution over S. The probability of w starting from μ under the guidance of π, denoted $\Pr_\mu^\pi(w)$, is equal to $\sum_{s \in S} \mu(s) \cdot \Pr_s^\pi(w, s)$, where $\Pr_s^\pi(w, \sigma) = 1$ if w is empty. If $w = aw'$,

$$\Pr_s^\pi(w, \sigma) = \sum_{s \xrightarrow{a} \mu} \pi(\sigma)(a, \mu) \cdot \sum_{t \in S} \mu(t) \cdot \Pr_t^\pi(w', \sigma \circ (a, t)).$$

Below follows the definition of trace distribution equivalence [10]:

Definition 7.4. *Let $\mathcal{A} = (S, Act, \rightarrow, L, \alpha)$ be a probabilistic automaton, and $\mu, \nu \in Dist(S)$. Then μ and ν are* trace distribution equivalent, *written as $\mu \simeq \nu$, if for each scheduler π of $\mathcal{A}$, there exists another scheduler π' such that $\Pr_\mu^\pi(w) = \Pr_\nu^{\pi'}(w)$ for each $w \in Act^*$ and vice versa.*

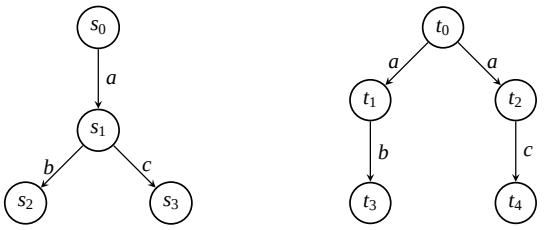

Figure 4: $\delta_{s_0} \simeq \delta_{t_0}$ but $\delta_{s_0} \not\sim_\dagger \delta_{t_0}$.

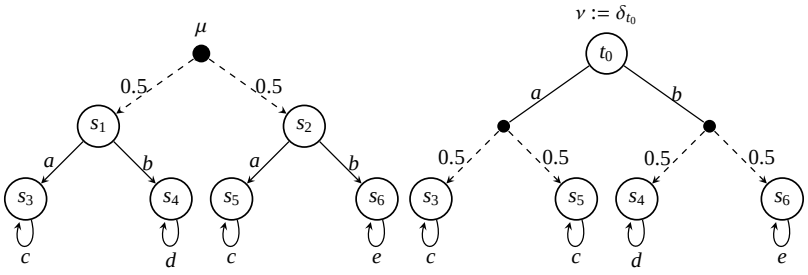

Figure 5: $\mu \sim_\S \nu$, but $\mu \neq \nu$.

Due to the existence of non-deterministic choices in a probabilistic automaton, we can have an alternative definition of trace distribution equivalence, called *a priori trace distribution equivalence*, by switching the order of the qualifiers of schedulers and traces, which resembles the definitions of *a priori bisimulation* in [24, 51].

Definition 7.5. *Let $\mathcal{A}, \mu$, and ν be as in Definition 7.4. Then μ and ν are* a priori trace distribution equivalent, *written as $\mu \simeq_{prio} \nu$, if for each scheduler π of $\mathcal{A}$ and $w \in Act^*$, there exists another scheduler π' of $\mathcal{A}$ such that $\Pr^\pi_\mu(w) \geq \Pr^{\pi'}_\nu(w)$ and vice versa.*

In Definition 7.4, we require $\Pr^\pi_\mu(w) \geq \Pr^{\pi'}_\nu(w)$ instead of $\Pr^\pi_\mu(w) = \Pr^{\pi'}_\nu(w)$, mainly to simplify the proofs in the sequel. However, due to the existence of combined transitions, these two definitions make no difference. Directly from the definitions, $\simeq$, $\simeq_{prio}$, and $\sim$ coincide on reactive automata. For general probabilistic automata, $\simeq_{prio}$ is strictly coarser than $\simeq$. Moreover, we have the following theorem:

Theorem 7.2. *1. $\simeq$ is incomparable to $\sim_\dagger$, $\sim_\S$, and $\sim$.*

2. $\sim \subsetneq \simeq_{prio}$.

PROOF. By Theorem 7.1, to prove clause 1 it suffices to show $\simeq \not\subseteq \sim$, $\sim_\S \not\subseteq \simeq$, and $\sim_\dagger \not\subseteq \simeq$.

1. $\simeq \not\subseteq \sim$. Let s_0 and t_0 be two states as in Fig. 4. It is easy to see that $\delta_{s_0} \simeq \delta_{t_0}$ but $\delta_{s_0} \not\sim \delta_{t_0}$.

2. $\sim_\dagger \not\subseteq \simeq$. Let μ and ν be as in Fig. 3. We have shown in the proof of Theorem 7.1 that $\mu \sim_\dagger \nu$. However, $\mu \not\simeq \nu$. For instance, there exists a scheduler of μ enabling us to see traces "ac", "ad", or "bd" each with probability $\frac{1}{3}$, which is not possible for ν.

3. $\sim_\S \not\subseteq \simeq$. Let μ and ν be as in Fig. 5. Let $R = \{(\mu, \nu), (\nu, \mu)\} \cup ID$. By Definition 7.1, R is a §-bisimulation. Therefore $\mu \sim_\S \nu$. However $\mu \not\simeq \nu$. For instance, there exists a scheduler such that from μ we will see "*be*" with probability $\frac{1}{2}$ while never see "*bd*", but starting from ν, the probabilities of "*be*" and "*bd*" are always the same.

4. $\sim \subseteq \simeq_{prio}$. By contraposition. Assume there exists μ and ν such that $\mu \sim \nu$ but $\mu \not\simeq_{prio} \nu$, i.e., there exists a scheduler π and a trace $w \in Act^*$ such that for all schedulers π', $\Pr^\pi_\mu(w) < \Pr^{\pi'}_\nu(w)$. Let $w = a_0 a_1 \ldots a_n$. Let μ_0 and ν_0 be the corresponding distributions of μ and ν after adding extra transitions to the deadlock state $\perp$. Then

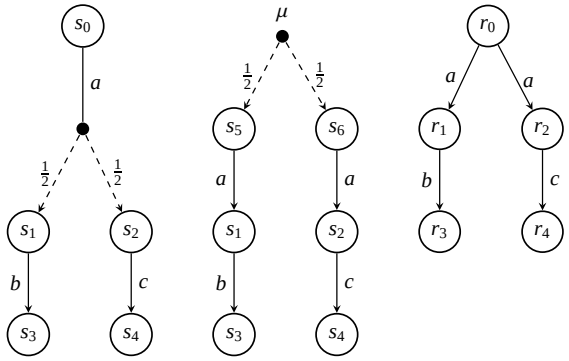

Figure 6: $\delta_{s_0} \sim \mu$, but $(\delta_{s_0} \parallel_{\mathbb{A}} \delta_{r_0}) \nsim (\mu \parallel_{\mathbb{A}} \delta_{r_0})$, where $\mathbb{A} = \{a, b, c\}$.

in the input-enabled extended automaton, let the chain of transitions $\mu_0 \xrightarrow{a_0} \mu_1 \xrightarrow{a_1} \ldots \xrightarrow{a_n} \mu_n$ mimic π as follows: For each $0 \leq i < n$ and $s \in supp(\mu_i)$, all transitions of s chosen by π with labels different from a_i are switched to transitions with labels a_i. After doing so, the probability of seeing w will not be lowered. By assumption, whenever $\nu_0 \xrightarrow{a_0} \nu_1 \xrightarrow{a_1} \ldots \xrightarrow{a_n} \nu_n$, it holds $\mu_n(\perp) < \nu_n(\perp)$, which contradicts that $\mu \sim \nu$.

5. $\simeq_{prio} \nsubseteq \sim$. The counterexample in Fig. 4 also applies here, as $\delta_{s_0} \simeq_{prio} \delta_{t_0}$ but $\delta_{s_0} \nsim \delta_{t_0}$. □

7.2. Compositionality

In this subsection, we discuss the compositionality of all mentioned bisimulation relations. Let $\mathcal{A}_i = (S_i, Act_i, \rightarrow_i, L_i, \alpha_i)$ be two probabilistic automata with $i \in \{0, 1\}$ and $\mathbb{A} \subseteq Act_1 \cap Act_2$. For any $\mu_0 \in Dist(S_0)$ and $\mu_1 \in Dist(S_1)$, we denote by $\mu_0 \parallel_{\mathbb{A}} \mu_1$ a distribution over $S_0 \times S_1$, the element of which is written as $s_0 \parallel_{\mathbb{A}} s_1$ where $s_0 \in S_0$ and $s_1 \in S_1$ for convenience, such that $(\mu_0 \parallel_{\mathbb{A}} \mu_1)(s_0 \parallel_{\mathbb{A}} s_1) = \mu_0(s_0) \cdot \mu_1(s_1)$. We recall the definition of parallel operator of probabilistic automata given in [10].

Definition 7.6. *Let $\mathcal{A}_i = (S_i, Act_i, \rightarrow_i, L_i, \alpha_i)$ be two probabilistic automata with $i \in \{0, 1\}$ and $\mathbb{A} \subseteq Act_1 \cap Act_2$. The parallel composition of $\mathcal{A}_1$ and $\mathcal{A}_2$ with respect to $\mathbb{A}$, denoted $\mathcal{A}_1 \parallel_{\mathbb{A}} \mathcal{A}_2$, is a probabilistic automaton $(S, Act_0 \cup Act_1, \rightarrow, L, \alpha_0 \parallel_{\mathbb{A}} \alpha_1)$ where*

- $S = S_0 \times S_1$,
- $s_0 \parallel_{\mathbb{A}} s_1 \xrightarrow{a} \mu_0 \parallel_{\mathbb{A}} \mu_1$ *with* $a \in \mathbb{A}$ *if* $\forall i \in \{0, 1\}$, $s_i \xrightarrow{a}_i \mu_i$,
- $s_0 \parallel_{\mathbb{A}} s_1 \xrightarrow{a} \mu_0 \parallel_{\mathbb{A}} \mu_1$ *with* $a \notin \mathbb{A}$ *if* $\exists i \in \{0, 1\}$, $s_i \xrightarrow{a}_i \mu_i$ *and* $\mu_{1-i} = \delta_{s_{1-i}}$,
- $L(s_0 \parallel_{\mathbb{A}} s_1) = L_0(s_0) \cup L_1(s_1)$.

We say $\sim$ is compositional if for any probabilistic automata $\mathcal{A}_0$ and $\mathcal{A}_1$, $\mathbb{A} \subseteq Act_0 \cap Act_1$, and any distributions $\mu_0, \mu_0' \in Dist(S_0)$ and $\mu_1 \in Dist(S_1)$, whenever $\mu_0 \sim \mu_0'$ in $\mathcal{A}_0$, we have $\mu_0 \parallel_{\mathbb{A}} \mu_1 \sim \mu_0' \parallel_{\mathbb{A}} \mu_1$ in $\mathcal{A}_0 \parallel_{\mathbb{A}} \mathcal{A}_1$. Similarly, we can define the notion of compositionality for the other relations. In the following, we show that all the three bisimulation relations mentioned in this section are unfortunately not compositional in general:

Theorem 7.3. *$\sim$, $\sim_{\S}$, and $\sim_{\dagger}$ are not compositional.*

PROOF. We only show the non-compositionality of $\sim$, since the other two can be proved in a similar way. Let s_0, μ, and r_0 be as in Fig. 6. It is easy to see that $\delta_{s_0} \sim \mu$. However, when composing δ_{s_0} and μ with δ_{r_0} by enforcing synchronization on $\mathbb{A} = \{a, b, c\}$, we have $(\delta_{s_0} \parallel_{\mathbb{A}} \delta_{r_0}) \nsim (\mu \parallel_{\mathbb{A}} \delta_{r_0})$. For instance, $\mu \parallel_{\mathbb{A}} \delta_{r_0}$ can reach the distribution $\frac{1}{2}\delta_{s_1 \parallel_{\mathbb{A}} r_1} + \frac{1}{2}\delta_{s_2 \parallel_{\mathbb{A}} r_2}$, from which transitions with label b or c are both available, each with probability $\frac{1}{2}$. However, this is not possible in $\delta_{s_0} \parallel_{\mathbb{A}} \delta_{r_0}$, where transitions with label b or c cannot be enabled at the same time. □

Although the three bisimulations are not compositional in general, we show that, by restricting to a subclass of schedulers, they are compositional. For this, we need to introduce some notations. Let $\hookrightarrow$ be a transition relation on

distributions defined as follows: In case μ is sequential, i.e., none of the states in its support has a Cartesian form, $\mu \overset{a}{\hookrightarrow} \mu'$ iff $\mu \overset{a}{\to} \mu'$ as in Definition 5.1. Otherwise if $\mu = \mu_0 \parallel_{\mathbb{A}} \mu_1$ for some $\mathbb{A} \subseteq Act$, then $\mu \overset{a}{\hookrightarrow} \mu'$ iff

- either $a \in \mathbb{A}$, and for all $i \in \{0,1\}$, there exists $\mu_i \overset{a}{\hookrightarrow} \mu_i'$ such that $\mu' = \mu_0' \parallel_{\mathbb{A}} \mu_1'$,
- or $a \notin \mathbb{A}$, and there exists $i \in \{0,1\}$ and $\mu_i \overset{a}{\hookrightarrow} \mu_i'$ such that $\mu' = \mu_0' \parallel_{\mathbb{A}} \mu_1'$ with $\mu_{1-i}' = \mu_{1-i}$.

Intuitively, $\hookrightarrow$ is subsumed by $\to$ such that transitions of a distribution $\mu_0 \parallel_{\mathbb{A}} \mu_1$ can be projected to transitions of μ_0 and μ_1. The definition of $\hookrightarrow$ can be generalized to distributions composed of more than 2 distributions. It is worthwhile to mention that the definition of $\hookrightarrow$ is not ad hoc. Actually, it coincides with transitions induced by *distributed schedulers* [52]. It has been argued by many authors, see e.g. [10, 53, 52], that schedulers defined in Definition 7.3 are too powerful in certain scenarios. For this, a subclass of schedulers, called *distributed schedulers*, was introduced in [52] to restrict the power of general schedulers. Instead of giving the formal definition of distributed schedulers, we illustrate the underlying idea by an example. We refer interested readers to [52] for more details.

Example 7.1. *Let s_0, r_0, μ be as in Fig. 6. We have shown in Theorem 7.3 that $\delta_{s_0} \parallel_{\mathbb{A}} \delta_{r_0} \nsim \mu \parallel_{\mathbb{A}} \delta_{r_0}$ with $\mathbb{A} = \{a, b, c\}$. The main reason is that in $\mu \parallel_{\mathbb{A}} \delta_{r_0}$, there exists a scheduler such that it chooses the left a-transition of r_0 when at $s_5 \parallel_{\mathbb{A}} r_0$, while it chooses the right a-transition of r_0 when at $s_6 \parallel_{\mathbb{A}} r_0$. This scheduler cannot be simulated by any scheduler of $\delta_{s_0} \parallel_{\mathbb{A}} \delta_{r_0}$. However, such a scheduler is not distributed, since it corresponds to two different transitions of r_0, thus cannot "distribute" its choices to states s_0 and r_0.*

Let $\mathcal{S}_D$ denote the set of all distributed schedulers. It is easy to check that distributed schedulers induce exactly transitions in $\hookrightarrow$. Even though $\sim$ is not compositional in general, we show that it is compositional if restricted to distributed schedulers, similarly for $\sim_{\S}$ and $\sim_{\dagger}$. Below we redefine the bisimulation relations with restricted to schedulers in $\mathcal{S}_D$, which is almost the same as Definition 5.2 except that all transitions under consideration must be induced by a distributed scheduler.

Definition 7.7. *Let $\mathcal{A} = (S, Act, \to, L, \alpha)$ be an input-enabled probabilistic automaton. A symmetric relation $R \subseteq Dist(S) \times Dist(S)$ is a (distribution-based) bisimulation with respect to $\mathcal{S}_D$ if $\mu R \nu$ implies that*

1. *$\mu(A) = \nu(A)$ for each $A \subseteq AP$, and*
2. *for each $a \in Act$, whenever $\mu \overset{a}{\hookrightarrow} \mu'$ then there exists $\nu \overset{a}{\hookrightarrow} \nu'$ such that $\mu' R \nu'$.*

We write $\mu \sim^{\mathcal{A}}_{\mathcal{S}_D} \nu$ if there is a bisimulation R with respect to $\mathcal{S}_D$ such that $\mu R \nu$.

In an analogous way, we can also define the restricted version of $\sim_{\S}$, $\sim_{\dagger}$, and $\sim_{\varepsilon}$, denoted $\sim_{(\S,\mathcal{S}_D)}$, $\sim_{(\dagger,\mathcal{S}_D)}$, and $\sim_{(\varepsilon,\mathcal{S}_D)}$ respectively. Below we show that by restricting to distributed schedulers, $\sim$, $\sim_{\S}$, and $\sim_{\dagger}$ are all compositional.

Theorem 7.4. *$\sim_{\mathcal{S}_D}$, $\sim_{(\S,\mathcal{S}_D)}$, and $\sim_{(\dagger,\mathcal{S}_D)}$ are compositional.*

PROOF. We only prove the compositionality of $\sim_{\mathcal{S}_D}$ here, as the proofs for the other cases are similar. Let $\mathcal{A}_i = (S_i, Act_i, \to_i, L_i, \alpha_i)$ be two probabilistic automata with $i \in \{0,1\}$ and $\mathbb{A} \subseteq Act_1 \cap Act_2$. Let $R = \{(\mu_0 \parallel_{\mathbb{A}} \mu_1, \nu_0 \parallel_{\mathbb{A}} \mu_1) \mid \mu_0 \sim_{\mathcal{S}_D} \nu_0\}$, where $\mu_0, \nu_0 \in Dist(S_0)$ and $\mu_1 \in Dist(S_1)$. It suffices to show that R is a bisimulation with respect to $\mathcal{S}_D$. Let $(\mu_0 \parallel_{\mathbb{A}} \mu_1) R (\nu_0 \parallel_{\mathbb{A}} \mu_1)$. Obviously, $(\mu_0 \parallel_{\mathbb{A}} \mu_1)(A) = (\nu_0 \parallel_{\mathbb{A}} \mu_1)(A)$ for each $A \subseteq AP$, since, say, $(\mu_0 \parallel_{\mathbb{A}} \mu_1)(A) = \sum_{B,B'.B \cup B' = A} \mu_0(B) \cdot \mu_1(B')$. Let $\mu_0 \parallel_{\mathbb{A}} \mu_1 \overset{a}{\hookrightarrow} \mu'$. We show that there exists $\nu_0 \parallel_{\mathbb{A}} \mu_1 \overset{a}{\hookrightarrow} \nu'$ such that $\mu' R \nu'$. We distinguish two cases:

1. $a \notin \mathbb{A}$: According to the definition of $\hookrightarrow$, either (*i*) $\mu_0 \overset{a}{\hookrightarrow} \mu_0'$ such that $\mu' \equiv \mu_0' \parallel_{\mathbb{A}} \mu_1$, or (*ii*) $\mu_1 \overset{a}{\hookrightarrow} \mu_1'$ such that $\mu' \equiv \mu_0 \parallel_{\mathbb{A}} \mu_1'$, We first consider case (*i*). Since $\mu_0 \sim_{\mathcal{S}_D} \nu_0$, there exists $\nu_0 \overset{a}{\hookrightarrow} \nu_0'$ such that $\mu_0' \sim_{\mathcal{S}_D} \nu_0'$. Therefore $\nu_0 \parallel_{\mathbb{A}} \mu_1 \overset{a}{\hookrightarrow} \nu_0' \parallel_{\mathbb{A}} \mu_1$. According to the definition of R, we have $\mu' \equiv (\mu_0' \parallel_{\mathbb{A}} \mu_1) R (\nu_0' \parallel_{\mathbb{A}} \mu_1) \equiv \nu'$ as desired. The proof of case (*ii*) is similar and omitted here.
2. $a \in \mathbb{A}$: It must be the case that $\mu_0 \overset{a}{\hookrightarrow} \mu_0'$ and $\mu_1 \overset{a}{\hookrightarrow} \mu_1'$ such that $\mu' \equiv \mu_0' \parallel_{\mathbb{A}} \mu_1'$. Since $\mu_0 \sim_{\mathcal{S}_D} \nu_0$, there exists $\nu_0 \overset{a}{\hookrightarrow} \nu_0'$ such that $\mu_0' \sim_{\mathcal{S}_D} \nu_0'$. Hence there exists $\nu_0 \parallel_{\mathbb{A}} \mu_1 \overset{a}{\hookrightarrow} \nu_0' \parallel_{\mathbb{A}} \mu_1'$ such that $\mu' \equiv (\mu_0' \parallel_{\mathbb{A}} \mu_1') R (\nu_0' \parallel_{\mathbb{A}} \mu_1') \equiv \nu'$. □

Example 7.2. *Let s_0, r_0, and μ be as in Example 7.1. Since s_0 and μ are sequential, $\sim_{\mathcal{S}_D}$ degenerates to $\sim$, and $\delta_{s_0} \sim_{\mathcal{S}_D} \mu$. We can also show that $\delta_{s_0} \parallel_{\mathbb{A}} \delta_{r_0} \sim_{\mathcal{S}_D} \mu \parallel_{\mathbb{A}} \delta_{r_0}$ with $\mathbb{A} = \{a, b, c\}$. Intuitively, by restricting to distributed schedulers, r_0 cannot choose different transitions when at $s_5 \parallel_{\mathbb{A}} r_0$ or $s_6 \parallel_{\mathbb{A}} r_0$, thus transitions with label b or c cannot be enabled at the same time.*

Since the bisimilarity $\sim$ can be seen as a special case of approximate bisimulation with $\varepsilon = 0$, approximate bisimulation is in general not compositional either. However, by restricting to distributed schedulers, the compositionality also holds for (discounted) approximate bisimulations.

Theorem 7.5. *$\sim_{(\varepsilon,\mathcal{S}_D)}$ is compositional for any $\gamma \in (0, 1]$.*

PROOF. Let $\mathcal{A}_i = (S_i, Act_i, \rightarrow_i, L_i, \alpha_i)$ be two probabilistic automata with $i \in \{0, 1\}$ and $\mathbb{A} \subseteq Act_1 \cap Act_2$. Let $\{R_{(\varepsilon,\mathcal{S}_D)} \mid \varepsilon \geq 0\}$, where

$$R_{(\varepsilon,\mathcal{S}_D)} = \{(\mu_0 \parallel_{\mathbb{A}} \mu_1, \nu_0 \parallel_{\mathbb{A}} \mu_1) \mid \mu_0 \sim_{(\varepsilon,\mathcal{S}_D)} \nu_0\},$$

be a family of relations on $Dist(S_0) \times Dist(S_1)$. It suffices to show that each $R_{(\varepsilon,\mathcal{S}_D)}$ is a (discounted) approximate bisimulation with respect to $\mathcal{S}_D$. Note that $d_{AP}(\mu_0 \parallel_{\mathbb{A}} \mu_1, \nu_0 \parallel_{\mathbb{A}} \mu_1) = d_{AP}(\mu_0, \nu_0)$. The remaining proof is analogous to the proof of Theorem 7.4 and omitted here. □

A direct consequence of Theorem 7.5 is that the bisimulation distance is non-expansive under parallel operators, if restricted to distributed schedulers, i.e., $D_b(\mu_0 \parallel_{\mathbb{A}} \mu_1, \nu_0 \parallel_{\mathbb{A}} \mu_1) \leq D_b(\mu_0, \nu_0)$ for any μ_0, ν_0, and μ_1.

8. Decidability and Complexity

It has been proved in [21, Lem. 1] that every linear bisimulation R corresponds to a bisimulation matrix E of size $n \times m$ with $n = |S|$ and $1 \leq m \leq n$. Two distributions μ and ν are related by R iff $(\mu - \nu)E = 0$, where distributions are seen as vectors. Furthermore, by making use of the linear structure, a decision algorithm was presented in [21] for $\sim_{\dagger}$. It was also mentioned that, with slight changes, this algorithm can be applied to deal with both $\sim$ and $\sim_{\S}$. Interested readers can refer to [21] for details about the algorithm. However, we show in this section that the problem of deciding *approximate* bisimulation is more difficult: it is in fact undecidable when no discounting is permitted, while the discounted version is decidable but NP-hard.

In the remaining part of this section, we shall focus on approximate bisimulations with and without discounting. We first recall the following undecidable problem [54].

Theorem 8.1. *Let $\mathcal{A}$ be a reactive automaton and $\varepsilon \in (0, 1)$. The following problem is undecidable: Whether $\mathcal{A}$ is ε-empty, i.e., whether there exists $w \in Act^*$ such that $\mathcal{A}(w) > \varepsilon$.*

By making use of the following reduction, we show that approximate bisimulation without discounting is undecidable.

Lemma 8.1. *Let $\mathcal{A}$ be a reactive automaton with initial distribution α and $\varepsilon \in (0, 1)$. Let s be a state such that $L(s) = \emptyset$ and $s \xrightarrow{a} \delta_s$ for all $a \in Act$. Then $\mathcal{A}$ is ε-empty iff $\alpha \sim_{\varepsilon} \delta_s$.*

PROOF.

1. Suppose $\mathcal{A}$ is ε-empty. We show that $\alpha \sim_{\varepsilon} \delta_s$. Assume $\alpha \nsim_{\varepsilon} \delta_s$. By construction, it must be the case that from α a distribution μ is reached in finite steps such that $\mu(F) > \varepsilon$ with $F \subseteq S$ being the set of accepting states, which contradicts that $\mathcal{A}$ is ε-empty.

2. Suppose $\alpha \sim_{\varepsilon} \delta_s$. We show that $\mathcal{A}$ must be ε-empty. By contraposition, suppose there exists $w \in Act^*$ such that $\mathcal{A}(w) > \varepsilon$. Let $w = a_0 a_1 \ldots a_n$. This means that there exits $\alpha \xrightarrow{a_0} \mu_0 \xrightarrow{a_1} \ldots \xrightarrow{a_n} \mu_n$ such that $\mu_n(F) > \varepsilon$. Since δ_s can only reach itself, we have $\alpha \nsim_{\varepsilon} \delta_s$, a contradiction. □

Directly from Theorem 8.1 and Lemma 8.1, we reach a proposition as below showing the undecidability of approximate bisimulation without discounting.

Proposition 8.1. *The following problem is undecidable: Given $\varepsilon \in (0,1)$ and $\mu, \nu \in Dist(S)$, decide whether $\mu \sim_\varepsilon \nu$ without discounting.*

For discounted approximate bisimilarity, the problem turns out to be decidable. Instead of presenting the algorithm formally in this paper, we only sketch how the algorithm works. Intuitively, in the definition of discounted approximate bisimulation, the distance ε is discounted with γ at each step. Since γ is strictly less than 1, ε will for sure become larger than or equal to 1 in finite steps, in which case $\mu \sim_\varepsilon \nu$ for any μ and ν. This enables us to identify a finite set of pivotal distributions, which contains enough information for deciding discounted approximate bisimulation in a probabilistic automaton. However, we also note that the algorithms presented in [28, 29] for computing *state-based* approximate bisimilarity cannot be applied here. Even though we can identify a finite set of pivotal distributions, for each pivotal distribution there are infinitely many distributions approximately bisimilar with it. Therefore in the algorithm these infinite sets of distributions have to be represented symbolically, which makes the whole algorithm very involved. Actually we can show that deciding discounted approximate bisimulation is NP-hard.

Theorem 8.2. *The following decision problem is* NP*-hard: Given $\varepsilon, \gamma \in (0,1)$ and $\mu, \nu \in Dist(S)$, decide whether $\mu \sim_\varepsilon \nu$ with discounting factor γ.*

PROOF. Firstly, we recall the following NP-hard problem from, say, [55]: Given an undirected graph $G = (V, E)$ and $k \leq |V|$, decide whether there exists a clique in G with size larger than k. Note a clique is a sub-graph where every two vertexes are connected. We shall reduce the clique checking problem to the problem of deciding discounted approximate bisimulation. Our reduction is almost the same as [56], where the clique checking problem was reduced to the problem of deciding the consensus string in a hidden Markov chain. We sketch the construction as below: Fix an order over vertexes in $V = \{a_1, \ldots, a_n\}$ with $n = |V|$. Let $\mathcal{A}_G = (S, Act, \to, L, \alpha)$ be a reactive probabilistic automaton such that

1. $S = \cup_{1\leq i\leq n} S_i \cup \{s, t, r\}$ where for each i, $S_i = \{s_i^j\}_{1\leq j\leq n}$. To simplify the presentation, we let $s_i^{n+1} = t$ for each i in the sequel;
2. $Act = V \cup \{\tau\}$ and $\alpha = \delta_s$;
3. The transition relation $\to$ is defined as follows:
   (a) $s \xrightarrow{\tau} \mu$ such that for $1 \leq i \leq n$, $\mu(s_i^1) = \frac{\lambda_i}{\lambda}$, where $\lambda = \sum_{1\leq i\leq n} \lambda_i$ with $\lambda_i = 2^{deg(a_i)}$ and $deg(a_i)$ being the degree of vertex a_i;
   (b) for each $1 \leq i, j \leq n$, we distinguish several cases: If $j = i$, then $s_i^j \xrightarrow{a_i} s_i^{j+1}$; If $j \neq i$ and $(a_i, a_j) \notin E$, then $s_i^j \xrightarrow{\tau} s_i^{j+1}$; If $j \neq i$ and $(a_i, a_j) \in E$, then both $s_i^j \xrightarrow{a_j} \mu$ and $s_i^j \xrightarrow{\tau} \mu$ where $\mu(s_i^{j+1}) = \mu(r) = \frac{1}{2}$;
   (c) All other transitions not stated in the above items have the form $u \xrightarrow{a} \delta_r$ for $u \in S$ and $a \in Act$;
4. $L(t) = AP$ and $L(u) = \emptyset$ for each $u \neq t$. That is, t is the *only* accepting state.

To help understanding the construction, we present in Fig. 8 the probabilistic automaton corresponding to the undirected graph depicted in Fig. 7. For simplicity, the state r and all transitions leading to it are omitted. The order over vertexes is defined by $a \leq b \leq c \leq d$, and the four branchings of s after performing action τ correspond to, from left to right, d, c, b, a, respectively. The example is taken from [56].

By construction, all paths in $\mathcal{A}_G$ able to reach t are of length $n + 1$ and have the same probability $\frac{1}{\lambda}$. The size of the maximal clique in G is k iff there exists $\delta_s \xrightarrow{\tau} \xrightarrow{b_1} \cdots \xrightarrow{b_n} \mu$ such that $\mu(t) = \frac{k}{\lambda}$ and $|\{b_i\}_{1\leq i\leq n} \setminus \{\tau\}| = k$. Moreover, the set $\{b_i\}_{1\leq i\leq n} \setminus \{\tau\}$ constitutes the maximal clique in G. Note that μ is also the distribution reachable from δ_s where the probability of t is maximal. Since such μ can only be reached from δ_s after performing $n + 1$ transitions, we have $\delta_r \sim_\varepsilon \delta_s$ for any $\varepsilon \geq \gamma^{n+1} \cdot \frac{k}{\lambda}$. Therefore, for any given $\gamma \in (0,1)$, the size of the maximal clique of G is k iff $\delta_r \sim_{(\gamma^{n+1}\cdot\frac{k}{\lambda})} \delta_s$ but $\delta_r \not\sim_{(\gamma^{n+1}\cdot\frac{k-1}{\lambda})} \delta_s$. □

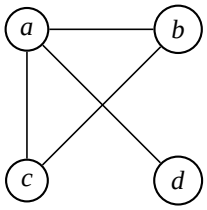

Figure 7: An undirected graph with $\{a, b, c\}$ being the maximal clique.

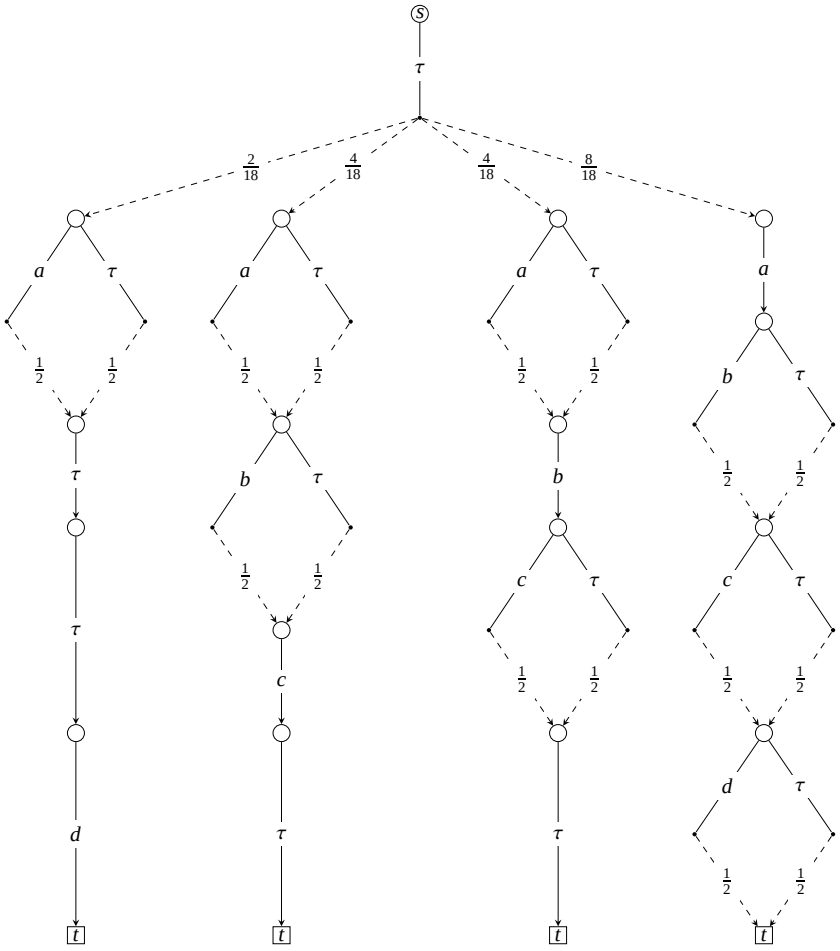

Figure 8: A reactive PA corresponding to the graph in Fig. 7.

9. Discussion and Future Work

In this paper, we considered Segala's automata, and proposed a novel notion of bisimulation by joining the existing notions of equivalence and bisimilarities. Our relations are defined over distributions. We have compared our bisimulation to some existing distribution-based bisimulations and discussed their compositionality and relations to trace equivalences. We have demonstrated the utility of our definition by studying distribution-based bisimulation metrics, which have been extensively studied for MDPs in state-based case. The decidability and complexity of deciding approximate bisimulations with or without discounting were also discussed.

State-based bisimulation has proven to be a powerful state space reduction technique in model checking. As future

work we would like to study how distribution-based bisimulations can be used to accelerate probabilistic model checking. One may combine it with state-based bisimulation which has efficient decision procedure, or component-based verification technique. As another direction of future work we would like to investigate weaker preorder relations such as simulations between distributions.

References

References

[1] M. Rabin, Probabilistic automata, Information and Control 6 (3) (1963) 230–245.

[2] W. Tzeng, A polynomial-time algorithm for the equivalence of probabilistic automata, SIAM Journal on Computing 21 (2) (1992) 216–227.

[3] S. Kiefer, A. S. Murawski, J. Ouaknine, B. Wachter, J. Worrell, Language Equivalence for Probabilistic Automata, in: CAV, Vol. 6806 of Lecture Notes in Computer Science, Springer, 2011, pp. 526–540.

[4] S. Kiefer, A. S. Murawski, J. Ouaknine, B. Wachter, J. Worrell, On the Complexity of the Equivalence Problem for Probabilistic Automata, in: FoSSaCS, Vol. 7213 of Lecture Notes in Computer Science, Springer, 2012, pp. 467–481.

[5] R. Bellman, Dynamic Programming, Princeton University Press, 1957.

[6] A. Bianco, L. de Alfaro, Model checking of probabalistic and nondeterministic systems, in: FSTTCS, Vol. 1026 of LNCS, Springer, 1995, pp. 499–513.

[7] M. Z. Kwiatkowska, G. Norman, D. Parker, PRISM 4.0: Verification of Probabilistic Real-Time Systems, in: CAV, Vol. 6806 of Lecture Notes in Computer Science, Springer, 2011, pp. 585–591.

[8] J.-P. Katoen, I. S. Zapreev, E. M. Hahn, H. Hermanns, D. N. Jansen, The ins and outs of the probabilistic model checker MRMC, Perform. Eval. 68 (2) (2011) 90–104.

[9] in: FM, Vol. 8442 of Lecture Notes in Computer Science, 2014.

[10] R. Segala, Modeling and Verification of Randomized Distributed Realtime Systems, Ph.D. thesis, MIT (1995).

[11] S. Cattani, R. Segala, Decision Algorithms for Probabilistic Bisimulation, in: CONCUR, Vol. 2421 of Lecture Notes in Computer Science, Springer, 2002, pp. 371–385.

[12] C. Baier, B. Engelen, M. E. Majster-Cederbaum, Deciding Bisimilarity and Similarity for Probabilistic Processes, J. Comput. Syst. Sci. 60 (1) (2000) 187–231.

[13] H. Hermanns, A. Turrini, Deciding Probabilistic Automata Weak Bisimulation in Polynomial Time, in: FSTTCS, Vol. 18 of LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2012, pp. 435–447.

[14] A. Parma, R. Segala, Logical Characterizations of Bisimulations for Discrete Probabilistic Systems, in: FoSSaCS, Vol. 4423 of Lecture Notes in Computer Science, Springer, 2007, pp. 287–301.

[15] J. Desharnais, V. Gupta, R. Jagadeesan, P. Panangaden, Weak bisimulation is sound and complete for pCTL*, Inf. Comput. 208 (2) (2010) 203–219.

[16] H. Hermanns, A. Parma, R. Segala, B. Wachter, L. Zhang, Probabilistic Logical Characterization, Inf. Comput. 209 (2) (2011) 154–172.

[17] L. Doyen, T. A. Henzinger, J.-F. Raskin, Equivalence of Labeled Markov Chains, Int. J. Found. Comput. Sci. 19 (3) (2008) 549–563.

[18] C. Eisentraut, H. Hermanns, L. Zhang, On Probabilistic Automata in Continuous Time, in: LICS, IEEE Computer Society, 2010, pp. 342–351.

[19] M. Hennessy, Exploring probabilistic bisimulations, part I, Formal Asp. Comput. 24 (4-6) (2012) 749–768.

[20] C. Eisentraut, J. C. Godskesen, H. Hermanns, L. Song, L. Zhang, Late weak bisimulation for Markov automata, CoRR abs/1202.4116.

[21] in: CONCUR 2014, Vol. 8704 of Lecture Notes in Computer Science, 2014.

[22] A. Giacalone, C. Jou, S. Smolka, Algebraic reasoning for probabilistic concurrent systems, in: IFIP TC2 Working Conference on Programming Concepts and Methods, North-Holland, 1990, pp. 443–458.

[23] J. Desharnais, V. Gupta, R. Jagadeesan, P. Panangaden, Metrics for Labeled Markov Systems, in: CONCUR, Vol. 1664 of Lecture Notes in Computer Science, Springer, 1999, pp. 258–273.

[24] L. de Alfaro, R. Majumdar, V. Raman, M. Stoelinga, Game relations and metrics, in: LICS, IEEE Computer Society, 2007, pp. 99–108.

[25] J. Desharnais, V. Gupta, R. Jagadeesan, P. Panangaden, Metrics for labelled Markov processes, Theor. Comput. Sci. 318 (3) (2004) 323–354.

[26] J. Desharnais, F. Laviolette, M. Tracol, Approximate Analysis of Probabilistic Processes: Logic, Simulation and Games, in: QEST, IEEE Computer Society, 2008, pp. 264–273.

[27] N. Ferns, P. Panangaden, D. Precup, Bisimulation Metrics for Continuous Markov Decision Processes, SIAM J. Comput. 40 (6) (2011) 1662–1714.

[28] D. Chen, F. van Breugel, J. Worrell, On the Complexity of Computing Probabilistic Bisimilarity, in: FoSSaCS, Vol. 7213 of Lecture Notes in Computer Science, Springer, 2012, pp. 437–451.

[29] H. Fu, Computing Game Metrics on Markov Decision Processes, in: ICALP (2), Vol. 7392 of Lecture Notes in Computer Science, Springer, 2012, pp. 227–238.

[30] G. Bacci, K. G. Larsen, R. Mardare, On-the-Fly Exact Computation of Bisimilarity Distances, in: TACAS, Vol. 7795 of Lecture Notes in Computer Science, Springer, 2013, pp. 1–15.

[31] G. Comanici, P. Panangaden, D. Precup, On-the-Fly Algorithms for Bisimulation Metrics, in: QEST, IEEE Computer Society, 2012, pp. 94–103.

[32] K. Chatzikokolakis, D. Gebler, C. Palamidessi, L. Xu, Generalized bisimulation metrics, in: CONCUR, Vol. 8704 of Lecture Notes in Computer Science, 2014, pp. 32–46.

[33] in: FM, Vol. 8442 of Lecture Notes in Computer Science, 2014.

[34] R. Segala, N. A. Lynch, Probabilistic Simulations for Probabilistic Processes, Nord. J. Comput. 2 (2) (1995) 250–273.

[35] L. Song, L. Zhang, J. C. Godskesen, Bisimulations meet PCTL equivalences for probabilistic automata, in: CONCUR, 2011, pp. 108–123.

[36] K. Larsen, A. Skou, Bisimulation through Probabilistic Testing, Information and Computation 94 (1) (1991) 1–28.
[37] B. Jonsson, K. Larsen, Y. Wang, Probabilistic extensions of process algebras, in: Handbook of Process Algebra, Elsevier, 2001, pp. 685–710.
[38] P. R. D'Argenio, N. Wolovick, P. S. Terraf, P. Celayes, Nondeterministic Labeled Markov Processes: Bisimulations and Logical Characterization, in: QEST, 2009, pp. 11–20.
[39] M. Bernardo, R. De Nicola, M. Loreti, A uniform framework for modeling nondeterministic, probabilistic, stochastic, or mixed processes and their behavioral equivalences, Inf. Comput. 225 (2013) 29–82.
[40] M. Stoelinga, F. W. Vaandrager, Root contention in IEEE 1394, in: Proceedings of the 5th Int. AMAST Workshop on Formal Methods for Real-Time and Probabilistic Systems, Vol. 1601 of LNCS, Springer, 1999, pp. 53–74.
[41] L. Zhang, H. Hermanns, F. Eisenbrand, D. N. Jansen, Flow faster: efficient decision algorithms for probabilistic simulations, in: TACAS, Vol. 4424 of LNCS, Springer, 2007, pp. 155–169.
[42] S. Crafa, F. Ranzato, Bisimulation and simulation algorithms on probabilistic transition systems by abstract interpretation, Formal Methods in System Design 40 (3) (2012) 356–376.
[43] J. Schuster, M. Siegle, Markov automata: Deciding weak bisimulation by means of non-naïvely vanishing states, Inf. Comput. 237 (2014) 151–173.
[44] C. Eisentraut, H. Hermanns, J. Krämer, A. Turrini, L. Zhang, Deciding bisimilarities on distributions, in: QEST, 2013, pp. 72–88.
[45] Y. Deng, R. J. van Glabbeek, M. Hennessy, C. Morgan, Testing finitary probabilistic processes, in: CONCUR, Vol. 5710 of Lecture Notes in Computer Science, Springer, 2009, pp. 274–288.
[46] M. Ying, M. Wirsing, Approximate bisimilarity, in: Algebraic Methodology and Software Technology, Springer, 2000, pp. 309–322.
[47] M. Ying, Topology in Process Calculus: Approximate Correctness and Infinite Evolution of Concurrent Programs, Springer, 2001.
[48] M. Ying, Bisimulation indexes and their applications, Theoretical Computer Science 275 (2002) 1–68.
[49] F. van Breugel, B. Sharma, J. Worrell, Approximating a Behavioural Pseudometric Without Discount for Probabilistic Systems, in: FoSSaCS, Vol. 4423 of Lecture Notes in Computer Science, Springer, 2007, pp. 123–137.
[50] K. Chatterjee, L. de Alfaro, R. Majumdar, V. Raman, Algorithms for game metrics (full version), Logical Methods in Computer Science 6 (3).
[51] M. Tracol, J. Desharnais, A. Zhioua, Computing distances between probabilistic automata, in: QEST, Vol. 57 of EPTCS, 2011, pp. 148–162.
[52] S. Giro, P. R. D'Argenio, Quantitative model checking revisited: Neither decidable nor approximable, in: FORMATS, Vol. 4763 of Lecture Notes in Computer Science, Springer, 2007, pp. 179–194.
[53] L. De Alfaro, The verification of probabilistic systems under memoryless partial-information policies is hard, Tech. rep., DTIC Document (1999).
[54] A. Paz, Introduction to Probabilistic Automata (Computer Science and Applied Mathematics), Academic Press, Inc., Orlando, FL, USA, 1971.
[55] M. R. Garey, D. S. Johnson, Computers and Intractability: A Guide to the Theory of NP-Completeness, W. H. Freeman & Co., New York, NY, USA, 1990.
[56] R. B. Lyngsø, C. N. S. Pedersen, The consensus string problem and the complexity of comparing hidden Markov models, J. Comput. Syst. Sci. 65 (3) (2002) 545–569.